

Bourreau, Marc; Gaudin, Germain

Working Paper

Streaming Platform and Strategic Recommendation Bias

CESifo Working Paper, No. 7390

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Bourreau, Marc; Gaudin, Germain (2018) : Streaming Platform and Strategic Recommendation Bias, CESifo Working Paper, No. 7390, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/191415>

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.

Streaming Platform and Strategic Recommendation Bias

Marc Bourreau, Germain Gaudin

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editors: Clemens Fuest, Oliver Falck, Jasmin Gröschl

www.cesifo-group.org/wp

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: www.CESifo-group.org/wp

Streaming Platform and Strategic Recommendation Bias

Abstract

We consider a streaming platform which carries content from various upstream content providers. Participating customers face personalized recommendations from the platform and consume a mix of content originating from each provider. We analyze when the platform uses its personalized recommendation system to steer consumers from one content provider to another. We establish the conditions under which the recommendation system allows the platform to credibly threaten upstream providers to steer consumers away from their content in order to reduce their market power. We find that the streaming platform can increase its profit by reducing the royalty rate it pays to content providers through the use of a recommendation system which is strategically biased in favor of the cheaper content.

JEL-Codes: D400, L100, L500.

Keywords: streaming platform, recommendation system, personalization, bias.

Marc Bourreau
Department of Economics and
Social Sciences, Telecom ParisTech
Paris / France
marc.bourreau@telecom-paristech.fr

Germain Gaudin
European Commission (Directorate-
General for Competition)
Brussels / Belgium
germain.gaudin@ec.europa.eu

October 9, 2018

We thank Emilio Calvano, Alexandre de Cornière, Axel Gautier, Doh-Shin Jeon, Igor Letina, André Marchand, Johannes Münster, Volker Nocke, Andrew Rhodes, Greg Taylor, Tommaso Valletti, and Joel Waldfogel for helpful comments and suggestions, as well as conference participants at the Media Workshop (Zurich, 2016), the PET Conference (Paris, 2017), the AFSE Conference (Paris, 2018), the EARIE Conference (Athens, 2018), and seminar participants in Cologne, Nice, at Telecom ParisTech, and at the Toulouse School of Economics. The views expressed in this article are personal, and do not necessarily represent those of the European Commission.

1 Introduction

Media streaming platforms, such as Spotify or Netflix, provide their customers with access to a broad range of content (music, movies, etc.), from a variety of providers. In order to bring consumers to subscribe to their service or to increase customer retention rates, streaming platforms typically set up sophisticated recommendation systems, which offer personalized recommendations to users. Personalized recommendations can be, for instance, based on users' past behavior (i.e., previous purchases or consumption), as well as on the information obtained through surveys or feedback (likes vs. dislikes, or ratings).¹ In turn, subscribing customers largely rely on platforms' recommendations to consume content. For instance, in 2016, Spotify announced that its personalized playlists ('Discover Weekly') were used by 40 million of its users (out of a total of 100 million users).² In a similar way, in 2013, Netflix estimated that 75 percent of its viewer activity was driven by recommendations.³

Recommendation systems can also serve a different, strategic purpose for streaming platforms. When a platform controls an integrated recommendation system, it can easily shift away from a situation where recommendations are used merely to increase customer retention or usage, and instead consider the overall profitability of the service that is recommended. In particular, when a platform pays royalties to content providers, it may have an incentive to bias its recommendations in order to steer consumers away from the most expensive content and towards the cheapest one.

For instance, in the music streaming industry, the online radio company Pandora revealed that it manipulates its recommendation algorithm in order to increase or decrease the frequency at which a music title is played based on the ownership of the sound recordings.⁴ In 2014, Pandora engaged into a special agreement with the indie-label coalition Merlin, whereby Merlin would accept reduced royalty rates in exchange for an increased performance of its titles. In practice, as put by Pandora, *"the Merlin agreement provides that as Pandora increases its performances of covered recordings – i.e., as Pandora "steers" toward Merlin-label recordings and away from competing recordings – its effective rate drops. [...] Pandora has precisely that ability to "steer" towards or away from the music of particular record companies."*⁵

¹See the comprehensive work edited by Ricci et al. (2011) for an overview of the design of recommendation systems.

²Source: Popper (2016).

³Source: Vanderbilt (2013). See also, e.g., Pathak et al. (2010), Oestreicher-Singer and Sundararajan (2012), and Aguiar and Waldfogel (2018) for some evidence of the impact of recommendations and playlists on consumption or usage.

⁴Pandora revealed that such "steering experiments" were part of its standard, business-oriented research investigations. See the testimony of Stephen McBride, Docket No. 14-CRB-0001-WR.

⁵See the Introductory memorandum to the written direct statement of Pandora Media, Inc., Docket No. 14-CRB-0001-WR.

In the movie streaming industry, Netflix and other movie streaming platforms have also been suspected to bias the personalized recommendations they make to users in order to favor their in-house productions.⁶ The practice of considering the overall profitability of a product or service when making recommendations to users is also used in online retailing.⁷

In practice, the pricing model of streaming platforms typically involves setting a (monthly) fee, against which a customer can access any content without additional charges.⁸ When the costs (e.g., royalties) of different content offered by the platform diverge, such pricing strategy prevents the platform to pass-on these cost differentials to consumers. In such a case, biasing recommendations can help platforms to overcome the problem arising from the fact that users do not internalize the differences in platforms' costs. Indeed, without biased recommendations, the resulting mix of content provided by a platform to a given consumer is likely to be sub-optimal from the streaming platform's perspective when it pays different royalty rates to content providers.

Moreover, to the extent that recommendation systems can steer users' streaming decisions, they can be used as a strategic tool through their impact on royalty rates paid by the platform to content providers.⁹ Streaming platforms' ability to use recommendation systems to steer subscribers away from (or towards) specific content, in return, can serve as a credible threat to reduce the market power of the upstream content providers.

In this paper, we examine the strategic use of a user-specific recommendation system by a monopoly streaming platform that carries content from two content providers. We study how the platform's ability to use its recommendation system can strategically affect the market power exercised by the upstream content providers.

We build up a model where a monopoly streaming platform offers content from two content providers. Consumers are heterogeneous in the optimal content-mix that they would like to consume, but they all derive the same utility if they obtain their optimal mix. On the supply side, the two providers set per-unit royalty fees to the platform for

⁶*House of Cards*, one of Netflix's most-successful in-house production, is said to be recommended to most of its subscribers, regardless of their past user-behavior. As presented in a New York Times article, "[g]iven that Netflix is in the business of recommending shows or movies, might its algorithms tilt in favor of the work it commissions as it goes deeper into original programming?" Source: Carr (2013). Also, one of the key drivers of Netflix's early success was a filter to its recommendation system screening for movies which were out-of-stock, thus reducing Netflix's costs by increasing the consumption of already acquired DVDs and lowering that of new releases, as explained by Shih and Kaufman (2014).

⁷For instance, after having purchased on Amazon, consumers may receive recommendation emails suggesting new products they may be interested in. In such case, as put by an Amazon spokesperson, "if a customer qualifies for both a Books mail and a Video Games mail, the email with a higher average revenue-per-mail-sent will win out." Source: Mangalindan (2012).

⁸See, e.g., Thomes (2013) for a model of streaming platforms, and Belleflamme (2016) for a recent survey of the related literature.

⁹When contracting with Merlin, the online radio Pandora explained that it would recommend Merlin artists over the non-affiliated ones in exchange for lower royalty rates. Source: Sydell (2014).

access to their content. Given the royalty rates, the platform sets a subscription fee but no usage fee, and decides on the design of its (personalized) recommendation system. The recommendation system is a technology which recommends a specific content-mix to each consumer, based on her type. If the platform recommends a content-mix different from a consumer's optimal content-mix, we say that there is a *recommendation bias*.

We begin our analysis with the case with exogenous per-unit royalty fees, under which we characterize the platform's optimal pricing strategy and recommendation bias. We find that when one content provider charges lower royalties than the other, the platform biases its recommendation to steer consumers towards the cheaper content. The platform faces a trade-off between setting a high subscription fee to maximize revenues and a high level of bias to minimize costs (i.e., royalties from content providers), because the platform needs to compensate (marginal) subscribers for bias via a lower subscription fee to ensure their participation. The larger the difference in royalty costs between the two providers, the larger the bias set by the monopoly platform at the equilibrium.

We then endogenize the per-unit royalty rates set by the upstream content providers. We show that the credibility of the platform's threat to steer its subscribers away from more costly content depends on how strong consumers feel about their favorite content-mix (in other words, consumers' flexibility to substitute one type of content for another). We find that if consumers are sufficiently insensitive to bias, the platform has the ability to steer and it can use its recommendation system strategically in order to obtain lower royalty rates from content providers.

Finally, we analyze various extensions of our baseline model, which provide additional insights. We show that when subscribers can search for their own content instead of following the platform's recommendation the platform's ability to steer remains if the consumers' search cost is not too low. We also study the platform's incentive to vertically integrate with one of the content providers. We find that vertical integration is another strategy that the platform can employ to mitigate the market power of the content providers, and that it is profitable when consumers are very sensitive to bias. We also investigate the possibility for the platform to use consumer information to set personalized subscription prices, and show that it allows the platform to earn higher profits. Finally, we show that if content providers offer some essential, "superstar" content, the platform has a higher ability to extract consumer surplus, and hence earns higher profits.

Our model can be applied to various digital platform markets such as online music streaming where consumers decide whether or not to pay a fixed fee in order to access a variety of content. The price paid by consumers does not depend on their content consumption, nor in volume, nor in type. Through personalized recommendations, the platform can gain some flexibility to tailor the consumption of each participating user.

Related literature. An extensive literature has highlighted the various features of recommendation systems which are useful in order to improve customer retention and to increase usage; see, e.g., Bakos (1997), Ansari, Essegai and Kohli (2000), Xiao and Benbasat (2007), Chung, Rust and Wedel (2009), as well as the comprehensive review edited by Ricci et al. (2011). However, these analyses mostly have a customer-oriented view of recommendation systems, and do not address how recommendations impact the overall profitability of the platform in the presence of (potential) asymmetries.¹⁰ Empirical analyses by Fleder and Hosanagar (2009), Brynjolfsson, Hu and Simester (2011), and Datta, Knox and Bronnenberg (2018) have provided some evidence on the impact of recommendation systems on consumer usage and the discovery of new products. However, they do not address directly the incentives of platforms with an integrated recommendation system to steer consumption through their recommendations.

Our paper also relates to the growing theoretical literature on search and recommendation bias, which so far has mainly focused on web-search engines. Various studies, such as Hagiu and Jullien (2011, 2014) and de Cornière and Taylor (2014), consider end-users that have free access to a search engine with an incentive to bias its recommendations towards its “preferred” items because of externalities resulting from, e.g., revenues from advertising markets.¹¹ The most closely related study to ours is by Calvano and Jullien (2018), who show that biases in recommendation systems can emerge in equilibrium in a duopoly setting. Their recommendation system, however, is neutral towards available items or content, and their focus is on platforms’ reputational concerns in a dynamic setting with asymmetric information.¹² Finally, in a related study building on a dynamic motivation for bias, Drugov and Jeon (2017) study the incentives of a vertically-integrated platform to bias recommendations towards its own content when consumers’ utility in the long-run is shaped by their short-run usage.

Our paper also relates to the literature on vertical markets, where retailers or intermediaries may attempt to steer end-consumers towards products that yield greater profit margins. This can be either because of high sales commissions paid by a particular producer as in Raskovich (2007), Armstrong and Zhou (2011), Inderst and Ottaviani (2012) or de Cornière and Taylor (2016), and can also occur when retailers set prices to final con-

¹⁰As a counterexample, Bodapati (2008) mentioned the necessary adjustments to an optimal recommendation system when recommended items or services differ in their profitability.

¹¹See also the related papers by Casadesus-Masanell and Halaburda (2014), and Burguet, Caminal and Ellman (2015).

¹²We do not consider any informational asymmetries. However, the underlying problem we emphasize resembles the standard moral hazard problem. The incentives of the platform are misaligned with users’ interests, as when financial advisers steer their clients towards investments that yield them higher commissions, or when medical professionals prescribe their patients with drugs produced by companies that give them “free lunches.”

sumers, as in Hunold and Muthers (2017).¹³ The main difference with our model, however, lies in the fact that in this literature consumers typically only purchase one single good or service, either from one manufacturer or the other, but never a mix of these. By contrast, in our setting the platform can shape the product or service that each participating user will consume by mixing the inputs obtained from various content providers. In media market, platform users are generally interested in consuming a mix of content from various sources, while only paying a single access price to the platform. Finally, our paper also broadly relates to the literature on price discrimination and product bundling because the platform offers consumers a personalized mix – or bundle – of products.¹⁴

The remainder of the paper is organized as follows. We begin with the setup of our model in Section 2, and describe our benchmark results in Section 3. We present the conditions under which the personalized recommendation system is biased towards one of the content provider’s content in Section 4, where we also endogenize royalty rates and discuss how the strategic use of recommendation system can benefit the platform. We present various extensions of our baseline model in Section 5. Finally, Section 6 presents several managerial implications of our results, and Section 7 concludes.

2 Model

A monopoly streaming platform charges a flat fee for access to content from two providers, providing personalized recommendations to its subscribers on their content-mix.

Upstream content providers. The two content providers, A and B , sell their content to the monopoly streaming platform with per-unit linear royalty rates denoted by r_A and r_B , respectively.¹⁵ For simplicity, we refer to the content provided by A and B , as A and B , respectively.

Downstream streaming platform. The streaming platform charges a fixed, uniform access price, P , to consumers, but no usage fee. The platform also makes recommendations

¹³See also Bardey et al. (2018) for a related analysis where multi-product firms can invest in order to uncover buyers’ preferences before offering one of their products.

¹⁴See, e.g., Adams and Yellen (1976), McAfee, McMillan and Whinston (1989) or, more recently, Jeon and Menicucci (2012).

¹⁵The marginal cost of providing access to content is often positive for streaming platforms. For instance, several streaming contracts have leaked over the past few years, such as that of January 2011 between Spotify and Sony Music Entertainment or that offered by Apple to independent music labels during Summer 2015, which showed that music streaming platforms may pay a per-play royalty to the content right owner. These contracts involve variable rates, and therefore our assumption of linear per-unit royalty rates can be seen as a reasonable approximation.

to participating consumers regarding their content-mix, which will be described below in detail.

Consumers. We consider that consumers are homogeneous with respect to their total consumption (streaming) of content, but heterogeneous with respect to their optimal content-mix. The idea is that A offers some content, B some other content, and consumers have heterogeneous preferences on how much to consume from each of them.

There is a unit mass of consumers uniformly distributed over the interval $[0, 1]$. Consumers' total consumption level is normalized to unity. The proportions of content A and content B in the optimal content-mix of a consumer of type x are x and $(1 - x)$, respectively.¹⁶

Consumers derive utility v from their optimal content-mix. The level of disutility a consumer incurs from a sub-optimal mix is quadratic in the distance between the consumer's optimal mix and her actual mix. That is, a consumer of type x who consumes a mix which consists of a share ρ of content A and, thus, $(1 - \rho)$ of content B , faces a disutility of $t(\rho - x)^2$, with $t > 0$.¹⁷

We assume that consumers and the platform are informed about consumers' optimal content-mix.¹⁸ If consumers do not participate (i.e., do not subscribe to the streaming service), they yield their outside option, which is assumed to be zero.

Recommendation system. The platform provides consumers with user-specific (personalized) recommendations. The platform's recommendation algorithm is defined as a function $\rho(\cdot)$, which takes a consumer's type x as an argument.

We say that the platform's recommendation to a given consumer of type x is *biased* if it differs from the consumer's optimal content-mix, which the platform would be able to recommend, that is, if $\rho(x) \neq x$.

Similar to Hagiu and Jullien (2011), and de Cornière and Taylor (2014), we assume that, if consumers decide to participate (i.e., subscribe), they always follow the recommended content-mix of the intermediary. That is, a consumer of type x consumes a share $\rho(x)$ of content A and $1 - \rho(x)$ of content B . In Section 5, we consider an extension of this baseline

¹⁶Anderson and Neven (1989) show that when consumers consume a product mix in a Hotelling location-price game, in equilibrium consumers attain their optimal product mix. This type of Hotelling model with combinable products has been later used, for example, by Gabszewicz, Laussel and Sonnac (2004) to model the demand for television channels.

¹⁷We assume that consuming a sub-optimal content-mix does not affect total consumption. The Pandora steering experiments reported that users' total radio consumption was rather inelastic in the recommendation bias: the experimental manipulations had a very minor, and (in most cases) statistically insignificant impact on the total number of hours consumers spent listening to Pandora. See the testimony of Stephen McBride, Docket No. 14-CRB-0001-WR.

¹⁸This assumption can be justified if the market is in a mature phase already, and consumers know their preferences well with regard to their consumption pattern.

model, where we allow consumers to ignore the platform's recommendation, and search for content from the two providers instead.

Timing. In the first stage of the game, content providers A and B set simultaneously their per-unit royalty rates, r_A and r_B . In the second stage, and given the per-unit royalty rates r_A and r_B , the platform sets its subscription price P and its recommendation algorithm $\rho(\cdot)$. Then, consumers decide whether to subscribe to the platform.

Finally, we make the technical assumption that content providers cannot set a royalty rate above v , the consumer's valuation when getting her favorite mix; that is, $r_i \leq v$, for $i = A, B$. A reason for such an upper bound is, e.g., that there is a competitive fringe of third-party content providers that supply content of a given type x at a royalty rate v .

Assumption 1. *Royalty rates are bounded by v .*

3 Benchmark: Non-Strategic Recommendation

In the benchmark scenario, we assume that the platform uses its recommendation system in order to boost demand as it focuses on customers' usage, but not to engage in some steering towards A or B . For example, the platform may purchase its recommendation system from a third-party provider, which follows its own objectives and does not internalize the platform's strategic objectives. The platform's recommendation system is then designed with the sole purpose to provide each customer with her optimal content-mix. The game is thus the same as the one presented in the previous section, except that the platform does not decide on the design of its recommendation system. We solve this benchmark game by reasoning backwards.

Second stage. In this benchmark, the platform's recommendation system is designed in order to provide each consumer with her optimal content-mix: that is, we have $\rho(x) = x$, for all $x \in [0, 1]$. Hence, there is no bias. Given that all consumers obtain their optimal content-mix, their net utility from subscribing to the platform's service is equal to $v - P$, for all $x \in [0, 1]$.

In the second stage of the game, the platform takes as given the royalties rates set by providers A and B . Without loss of generality, we assume that $r_B \geq r_A$, and define $\Delta \equiv r_B - r_A$.

The only strategic variable of the platform is its subscription price, P . It is clear in this case that the platform sets $P^* = v$, and that all consumers participate. Given that all consumers obtain their optimal content-mix, 50% of total consumption goes to content A ,

and 50% to content B . Hence, the platform's total costs are equal to $(r_A + r_B)/2$, and its profit is $v - (r_A + r_B)/2$.

First stage. In the first stage of the game, content providers set their royalty rates simultaneously. They anticipate that the platform recommends to each consumer her optimal content-mix. Hence, content A and content B are perfect complements to the platform. Content providers can thus capture the entire revenue v by setting royalty rates such that $r_A^* + r_B^* = 2v$. (If rates lead to a higher cost for the platform, it decides to shut down.) Given Assumption 1, the only equilibrium is symmetric, with $r_A^* = r_B^* = v$. Hence, the platform earns zero profit when the royalty rates are endogenized.¹⁹

In sum, in this benchmark scenario, total welfare is maximized as each consumer participates and obtains its optimal content-mix. The value created in the market is entirely captured by the content providers; the platform and consumers are left with zero surplus in equilibrium.

4 Strategic Recommendation

We now consider that the platform internalizes the impact of the recommendation system on its profit, e.g., because it is developed in-house. The platform can use its recommendation system not only to maximize consumers' valuation of the service, but also to steer consumption towards the cheaper content. As above, we solve the game by reasoning backwards.

4.1 Second stage

At the last stage of the game, the platform determines the profit-maximizing subscription price and recommendation system, given the royalty rates set by the content providers. Without loss of generality, we assume that content provider B offers more costly content to the platform than content provider A , i.e., $r_B \geq r_A$.

4.1.1 Consumer participation

When the platform has set its subscription price and designed its recommendation system, consumers decide whether to participate. They are aware of both their own tastes and the

¹⁹Note that when Assumption 1 is not satisfied, there exists a continuum of equilibria where r_A^* is equal to a share $\alpha \in [0, 1]$ of $2v$, and the other royalty rate, r_B^* , to the complement share, $2v(1 - \alpha)$. The only symmetric equilibrium is the one where $r_A^* = r_B^* = v$. In any case, the content providers jointly extract the entire surplus.

platform's recommendation system.²⁰ Therefore, a consumer of type x joins the platform if and only if:

$$v - t[\rho(x) - x]^2 - P \geq 0. \quad (1)$$

Since $r_B \geq r_A$, when the platform distorts recommendations it has an incentive to do so by favoring provider A . This implies that $\rho(x) \geq x$, for all $x \in [0, 1]$.

4.1.2 User-specific recommendations

We now study the profit-maximizing design of the recommendation system, for a given subscription price. Note that the personalized recommendation system does not exclude the possibility for the platform to offer a subset of consumers (possibly the entire set of consumers) their optimal content-mix.

The platform's margin for a consumer of type x who receives a recommendation $\rho(x)$ is denoted by $M(\rho(x), x) \equiv P - r_B + \rho(x)\Delta$, with $\Delta = r_B - r_A$. Since $\Delta \geq 0$, the platform's margin increases in $\rho(x)$; intuitively, the platform makes a higher margin if the consumer is steered towards the cheaper content. The platform thus chooses the highest recommendation $\rho(x)$ that satisfies the participation condition expressed in equation (1). Therefore, we have:

$$\rho(x) = \min\left\{1, x + \sqrt{(v - P)/t}\right\}. \quad (2)$$

Note that we have necessarily $P \leq v$, otherwise no consumer would subscribe to the service.

There exists $\hat{x}(P) \equiv 1 - \sqrt{(v - P)/t} \leq 1$ such that for $x \geq \hat{x}$ then $\rho(x) = 1$, whereas $\rho(x) = x + \sqrt{(v - P)/t}$ for $x < \hat{x}$.²¹ Consumers with a strong preference for the cheaper content ($x \geq \hat{x}$) are offered only content A , while the other consumers ($x < \hat{x}$) are offered a personalized mix of A and B . In both cases, consumers are offered a larger share of content A (the cheaper content) than predicted by their optimal content-mix.

We now plug the expression for $\rho(x)$ into the platform's margin for a consumer of type x . If $x \geq \hat{x}$, we have $\rho(x) = 1$ and the platform's margin for consumers of type x is equal to $P - r_A$. We have necessarily $P \geq r_A$, otherwise the platform would not cover its royalty costs and hence shut down. The platform's margin is therefore positive for all $x \geq \hat{x}$. Instead, if $x < \hat{x}$, we have $\rho(x) = x + \sqrt{(v - P)/t}$, and the platform's margin is positive if and only if $x \geq \hat{x}(P)$, where $\hat{x}(P) \equiv \hat{x}(P) - (P - r_A)/\Delta$. Again, we have necessarily $P \geq r_A$, and therefore $\hat{x} \leq \hat{x}$.

²⁰Our model thus suits better a situation where the market is in a mature phase already. This is precisely in this type of situation that the platform could find it profitable to deviate from a purely customer-centric view when designing its recommendation system, and also focus on overall profitability.

²¹Note that if x was distributed on a shorter interval than $[0, 1]$, it could be the case that no consumer would be recommended content from only one provider. See the extension in Section 5 where each content provider owns some "superstar" content demanded by all consumers.




Figure 1: Recommendation pattern with bias and exclusion.

Does the platform possess the ability to exclude consumers for which it would incur losses, for a given subscription price? The platform would indeed like to exclude the consumers of type $x < \hat{x}$ for which it earns a negative margin; that is, for which $P - r_B + \rho(x)\Delta < 0$, where $\rho(x)$ is given by equation (2) above. The platform can dissuade a consumer of type $x < \hat{x}$ to subscribe to its service by providing her sub-optimal recommendations, $\bar{\rho} > \rho(x)$, given the participation constraint mentioned in equation (1). Hence, whenever the platform has an incentive to exclude consumers of some type, it can do so.

Figure 1 illustrates the design of the recommendation system when there is recommendation bias and exclusion of some consumers. Consumers who have a relatively strong preference for the cheaper content ($x > \hat{x}$) are recommended only content A. Consumers who have a relatively strong preference for the costlier content ($x < \hat{x}$) are excluded with very biased recommendations. Finally, the platform offers a personalized content-mix to each consumer in between, which contains both A (in blue) and B (in red), and which ensures that the consumer subscribes.

One major difference of our model with respect to the existing literature on steering is that customers of streaming services do not purchase a discrete good, but, instead, consume a bundle of content. Therefore, a consumer's optimal consumption is generally

not to obtain 100% of content from the same provider. This allows the streaming platform to adjust its recommendations to consumer types in such a way that consumers of different types $x \in [\hat{x}, \bar{x}]$ are given different (personalized) recommendations. In equilibrium, the bias for these consumers remains constant across types. By contrast, when consumers purchase a discrete good the bias that they face typically depends on their type (see, e.g., Raskovich (2007)).

4.1.3 Subscription price

We now determine the subscription price set by the platform. To determine the profit-maximizing subscription price, we start by looking at the equilibrium with full consumer participation (i.e., with $\hat{x} \leq 0$). The platform sets P in order to maximize its profit, given by:

$$\Pi = \int_0^1 (P - r_A) dx - \Delta \int_0^{\dot{x}(P)} [1 - \rho(x)] dx. \quad (3)$$

Solving for the first-order condition, we find the profit-maximizing subscription price,

$$P^* = v - \frac{t\Delta^2}{(\Delta + 2t)^2}. \quad (4)$$

The second-order condition is satisfied if and only if $\Delta \geq 0$ and the solution is continuous at $\Delta = 0$, where $P^* = v$, which ensures that P^* is determinate for all $\Delta \geq 0$.

We now check whether the market is indeed covered at the profit-maximizing subscription price defined above. The platform's margin for a consumer of type x and a given subscription price P , denoted by $M(\rho(x), x)$, is increasing in x . Therefore, the market is fully covered if and only if $M(\rho(0), 0) \geq 0$ at $P = P^*$. This is true if and only if $v - r_B + \Delta^2(\Delta + t)/(\Delta + 2t)^2 \geq 0$, which is satisfied under Assumption 1.

We can further demonstrate that in equilibrium, the platform has always an incentive to cover the market and, hence, the equilibrium subscription price is equal to P^* , given above in equation (4).²² Replacing for P^* into $\dot{x}(P)$, we obtain:

$$\dot{x}(P^*) = \frac{2t}{\Delta + 2t}. \quad (5)$$

The platform thus provides a full bias $\rho(x) = 1$ to consumers with type higher than $2t/(\Delta + 2t)$. There is bias in equilibrium if and only if $\dot{x}(P^*) < 1$, which is true whenever $\Delta > 0$. In equilibrium, the content-mix recommended to each type x is given by $\rho(x) = \min\{1, x + \Delta/(\Delta + 2t)\}$. The equilibrium is characterized in Proposition 1 below.

²²See the proof of Proposition 1 in the Appendix.

Proposition 1. *Under Assumption 1, the platform sets the subscription price given by equation (4), and always supplies the entire market in equilibrium. The platform biases recommendations for each type $x < 1$ in equilibrium if and only if $\Delta > 0$, with $\rho(x) = 1$ for $x > 2t/(\Delta + 2t)$, and $\rho(x) = x + \Delta/(\Delta + 2t)$ otherwise.*

Proof. See Appendix. □

When it sets its subscription fee and designs its recommendation system, the platform trades off between setting a high fee to maximize revenues and implementing highly biased recommendations to minimize royalty payments. This is because the platform has to compensate marginal subscribers for bias via a lower subscription fee to ensure their participation. The larger the difference in royalty costs between the two content providers, the larger the bias set by the platform in equilibrium.

Because when $\Delta > 0$ bias is only partial for consumers of type $x < 2t/(\Delta + 2t)$, in equilibrium the platform offers consumer-specific (biased) recommendations. The platform's equilibrium profit is given by $v - r_A - \Delta t/(\Delta + 2t)$. Although it sets a uniform subscription price, the platform can bias its recommendation system in such a way that it extracts all available surplus from a large set of participating consumers, as shown in Figure 2. Hence, for any given pair of royalty rates, the platform (weakly) increases its profit by also considering its own costs when setting its recommendations to consumers. Interestingly, consumers who have a strong preference for the cheapest content obtain a strictly positive net surplus in equilibrium.²³

In order to understand the characterization of the equilibrium in our model, it is important to see that in most of the literature, the intermediary does not set final prices to consumers, but instead only benefits from (potentially asymmetric) marginal transfers from the product providers (see, e.g., the papers by Raskovich (2007), or Inderst and Ottaviani (2012)). Intermediaries set consumer prices in the model proposed by Hunold and Muthers (2017), but the optimal retail prices are independent from the level of service provided, as demand is separable in service and in retail prices. By contrast, in our model, the streaming platform sets prices to consumers and internalizes the revenue loss following a reduction of its subscription price, which is necessary to keep consumers on board when recommendations are biased. This direct effect of the recommendation bias on revenues is important in our model, as all consumers with types between 0 and \hat{x} are marginal because they obtain zero net surplus.

²³In Section 5, we discuss the robustness of this result by considering two different extensions: one in which the content providers offer “superstar” content, and another where the platform can price-discriminate.




Figure 2: Consumer net utility with bias and covered market.

4.2 First stage

We now determine the equilibrium royalties set by content providers in the first stage of the game. From the above analysis, the content providers' profits are given by:

$$\pi_A(r_A, r_B) = r_A \left[1 - \frac{2t^2}{(\Delta + 2t)^2} \right], \quad (6)$$

and:

$$\pi_B(r_A, r_B) = r_B \left[\frac{2t^2}{(\Delta + 2t)^2} \right]. \quad (7)$$

We can check for content providers' unilateral incentives to deviate from any state where $v \geq r_B \geq r_A$. We find that only two equilibria can arise, depending on the model parameters, which are both symmetric. This result is stated in the following proposition.

Proposition 2. *Under Assumption 1, the equilibrium royalty rates are given by (i) $r_A^* = r_B^* = t$ if $v \geq t$, and (ii) $r_A^* = r_B^* = v$ if $v < t$.*

Proof. See Appendix. □

At the equilibrium, given that $r_A^* = r_B^*$, there is no bias and each consumer is recommended her optimal mix of content A and B . Therefore, total surplus is maximized. Moreover, the strategic effect of recommendations on competition between content providers implies that the platform sets a subscription price equal to v , and consumers are left with zero surplus. Finally, the platform makes a positive profit in equilibrium when $v > t$, while the content providers capture the entire surplus when $t \geq v$.

4.3 Strategic effects of recommendation

We can now compare the equilibrium detailed above with that of our benchmark case, where the platform simply uses its recommendation system to maximize the consumer surplus it can extract through its subscription price but without taking into account the strategic effect of (potential) recommendation bias on royalty rates.

Proposition 3. *The strategic use of the recommendation system triggers (weakly) lower royalty rates in equilibrium. The platform strictly increases its profit due to this strategic effect when $v > t$.*

Proof. Omitted. □

When the platform can use its recommendation system strategically, it can steer consumers toward or away from the content of a given provider. The ability of the platform to steer depends on consumers' sensitivity to biased recommendations. In our framework, this is measured by the parameter t , which determines the magnitude of the consumer utility loss resulting from a biased recommendation. The lower t , the less sensitive to bias consumers are, and hence the stronger the ability of the platform to steer.

If consumers are sufficiently insensitive to bias (i.e., if $t < v$ in the proposition), the content providers face a platform with a strong ability to steer. Each one of them has then an incentive to set a lower royalty rate than its rival to induce the platform to steer consumers toward its content. This leads to an equilibrium with lower royalties compared to the benchmark with a non-strategic recommendation system, thus increasing the platform's profit.

Therefore, to the extent that consumers are not too sensitive to bias, the platform has an incentive to invest in a recommendation system, which confers the ability to reduce the market power of content providers.

5 Extensions

In this section, we consider different extensions to our baseline setting. First, we allow consumers to bypass the platform's recommendation and search for their optimal content-mix. Second, we investigate the impact of vertical integration between the platform and one content provider on the equilibrium outcome. Finally, we discuss the robustness of our finding that the platform cannot extract the full consumer surplus via two variations of our main setting: (i) when the platform can price-discriminate according to consumers' types; and (ii) when content providers offer "superstar" content that consumers view as essential in their content-mix.

5.1 Consumer search

In our baseline model, participating customers must follow the platform's recommendation. While this model applies well to online radio (e.g., Pandora) or curated playlists, streaming platforms such as Spotify also allow consumers to search for the music that they like and to create their own playlists.

In this extension, we thus assume that after having decided to subscribe to the platform's service, any participating customer can decide to bypass the platform's recommendation by incurring a search cost $s > 0$. This search cost corresponds to the cost of finding relevant content on the platform and to create a playlist. When they search for content, we assume that consumers attain their optimal content-mix and, hence, obtain the net utility $v - s - P$. Finally, we assume that $s < t$, otherwise consumers would never search.

In the benchmark scenario where the platform offers optimal recommendations $\rho(x) = x$, the equilibrium remains the same whether or not consumers have the ability to search. Indeed, consumers are always better-off following the platform's recommendation, as it allows them to obtain their optimal mix without incurring any search cost.

By contrast, in the model with strategic recommendations, allowing consumers to search for their optimal mix may change the equilibrium. First, recall that in the last stage of the game, the platform takes content providers' royalty rates as given. For any subscription price P , consumers obtain $v - P - s$ if they search on the platform, and $v - P - t[\rho(x) - x]^2$ if they follow the recommendation. Hence, the equilibrium given by Proposition 1 still holds when consumers can search, if $s \geq t\Delta^2/(\Delta + 2t)^2$. Indeed, in this case, search costs are so high that participating customers are better off following the (potentially biased) recommendations.

Instead, if $s < t\Delta^2/(\Delta + 2t)^2$, some consumers – including all marginal consumers with type $x \in [0, \hat{x}]$ – would prefer to search at the subscription price and the recommendations given in Proposition 1. The platform may thus have an incentive to alter its recommendation system and its subscription price in order to prevent consumer search. Indeed, a consumer of type x who searches derives the utility $v - P - s$, while the platform earns the margin $P - [xr_A + (1 - x)r_B]$ serving this consumer. The platform would be better off offering a recommendation $\rho(x) \geq x$ such that the disutility from bias for the consumer is equal to s , leaving the consumer with the same utility, while increasing the platform's margin to $P - [\rho(x)r_A + (1 - \rho(x))r_B]$. The platform thus sets the maximum bias possible which ensures that participating consumers do not search for their optimal content-mix. Therefore, for a given subscription price, the platform sets up the following recommendation system:

$$\rho(x) = \min \left\{ 1, x + \sqrt{s/t} \right\}. \quad (8)$$

All participating consumers with type $x \geq \bar{x} \equiv 1 - \sqrt{s/t}$, with $\bar{x} > \hat{x}$, are recommended 100% of content from provider A , whereas the others (with type $x < \bar{x}$) face a bias which is smaller than when search is not feasible. Finally, and given the optimal recommendations, the platform charges the subscription price $P^* = v - s$ in equilibrium under a covered market, in order to ensure consumer participation, which is larger than the equilibrium price obtained when consumers cannot search. The platform is able to set a larger subscription fee when consumers possess the ability to search, because its recommendations are less biased. Overall, the platform earns a lower profit when consumers are able to search for their optimal content-mix, because this represents a constraint on the platform's ability to steer consumption.²⁴

The first stage of the game may also be modified when search costs are small enough. For instance, if there were no search costs, consumers could disregard the platform's recommendations and freely consume their optimal mix. Anticipating this, content providers would demand the maximum royalties, $r_A = r_B = v$, in equilibrium. More generally, when $s > 0$ and $v \geq t$, content providers understand that the platform's steering ability through a biased recommendation system, following an unilateral upwards deviation in their royalty rate, is limited by the consumers' ability to search. Therefore, content providers may have an incentive to depart from the equilibrium detailed in Proposition 2 and to increase their royalty rate. This may jeopardize the existence of an equilibrium, as a large, unilateral price deviation downwards may become profitable when royalty rates are high. In order to guarantee the existence of a Nash equilibrium in pure strategies, we make the following assumption.

Assumption 2. *Consumers' search cost is not too small: $s \geq \hat{s} \equiv t(1 - \sqrt{t/v})^2$.*

Assumption 2 is the necessary and sufficient condition for the existence of an equilibrium in pure strategies, when $v \geq t$. We can now demonstrate the robustness of our result stated in Proposition 2 above to the case where consumers can decide to bypass the platform's recommendation.

Proposition 4. *When participating customers have the option to obtain their optimal content-mix at a positive search cost s , we have the following:*

- *If $v \geq t$, an equilibrium in pure strategies exists if Assumption 2 is satisfied, and the equilibrium royalty rates are given by $r_A^* = r_B^* = t$.*
- *If $v < t$, the equilibrium royalty rates are given by $r_A^* = r_B^* = v$.*

Proof. See Appendix. □

²⁴In this extended model with search, we find that the platform has no incentive to exclude some consumers or to set a full bias, as in the baseline model. The proof is available in the web appendix.

Proposition 4 thus shows that our results are robust to the possibility for consumers to search for content. The constraint exerted by the consumers' ability to search for their favorite content-mix does not alter our main result, stated in Proposition 3, that the strategic use of the recommendation system leads to (weakly) lower royalty rates in equilibrium. The platform thus reduces its costs and earns greater profits when using its recommendation system strategically.

5.2 Vertical integration

In the baseline model, we have considered that the content providers are independent from the platform. However, in reality, streaming platforms may exhibit some degree of vertical integration with content providers. For example, one of the leading music streaming platforms, Spotify, has been recently trying to bypass record labels by licensing music directly from some artists, which is tantamount to vertical integration.²⁵ Video streaming platforms, such as Netflix or Amazon Video, also propose a large share of content produced in-house.

In this extension, we therefore investigate an alternative setting where the platform is vertically integrated with one content provider, A . It can thus obtain content from A at the implicit royalty rate $r_A = 0$. In the first stage of the game, only content provider B sets its royalty rate. In this situation, the outcome of the game in the second-stage is still determined by Proposition 1, with $\Delta = r_B$.

In the first stage of the game, content provider B 's profit changes as follows after a small deviation in r_B , at $r_A = 0$:

$$\left. \frac{\partial \pi_B}{\partial r_B} \right|_{r_A=0} = \frac{-2t^2 (r_B - 2t)}{(r_B + 2t)^3}. \quad (9)$$

Therefore, content provider B 's optimal royalty rate is $r_B = \min\{2t, v\}$.²⁶ Since $r_A = 0$, royalty rates differ, which implies that the platform sets up a biased recommendation system in equilibrium.

Proposition 5. *When the platform is integrated with a content provider (here, A), recommendations are biased in equilibrium. In addition, in equilibrium:*

- If $v \geq 2t$, B 's royalty rate is $r_B^* = 2t$, the recommendation system is given by $\rho(x) = \min\{1, x + 1/2\}$, and the subscription price equals $P^* = v - t/4$.
- If $v < 2t$, B 's royalty rate is $r_B^* = v$, the recommendation system is given by $\rho(x) = \min\{1, x + v/(v + 2t)\}$, and the subscription price equals $P^* = v - tv^2/(v + 2t)^2$.

²⁵See Nicolaou (2018).

²⁶The second-order condition is satisfied at $r_B = 2t$.

Proof. See Appendix. □

When $v \geq 2t$, content provider B sets the royalty rate $r_B = 2t$. This means that the royalty rate for content B is set at a higher level than in the equilibrium of our baseline model. In this case, the subscription price equals $P^* = v - t/4$. The recommendation system is given by $\rho(x) = \min\{1, x + 1/2\}$. The asymmetry at the upstream level implies that there is some bias in equilibrium, and intuitively, the platform biases its recommendations towards in-house content. Besides, the platform earns a profit of $v - t/2$ in equilibrium, which is equal to the sum of its own profit plus that of content provider A in the baseline model. Content provider B captures a profit of $t/4$ when the platform is vertically integrated, whereas it obtained $t/2$ in the baseline model. Consumers, who were left with zero surplus under the baseline model, retain a surplus of $t/12$ under integration.

When $v < 2t$, the royalty rate r_B is set at its maximum, v , which yields the subscription price $P^* = v - tv^2/(v + 2t)^2$ and the recommendation system $\rho(x) = \min\{1, x + v/(v + 2t)\}$. The integrated platform earns $v - tv/(v + 2t)$, which is greater than the sum of profits from the platform itself and content provider A in the equilibrium of the baseline model. By contrast, the profit of content provider B is always lower under integration. Consumers are left with a positive surplus of $(2t/3)v^3/(v + 2t)^3$.

Consider now that the platform and content provider A have to incur a very small, but nonzero, fixed cost in order to merge. They then decide to merge if the sum of their profits increases strictly with the merger.

Proposition 6. *If there is a very small, but nonzero, fixed cost to vertical integration, the platform and content provider A have an incentive to merge if and only if $v < 2t$.*

Proof. Immediate from the analysis above. □

A merger between the platform and a content provider is profitable if consumers are sufficiently adverse to bias (i.e., t is high), in which case the content providers have a strong market power if they are independent from the platform. The platform can mitigate this market power by merging with one provider. The merger hurts the independent content provider, B , but benefits the merged entity and consumers.

5.3 Platform's profit and consumer surplus

When royalty rates are exogenously set, we have demonstrated above that the platform's strategic recommendation system increases consumer surplus for a given pair of royalty rates $r_B > r_A \geq 0$, compared to the benchmark scenario where $\rho(x) = x$, $\forall x$. (Recall that in this benchmark scenario consumers are left with zero surplus.) In this extension, we investigate the robustness of this finding, and we show that the platform may earn greater

profits by capturing (part of) the surplus left to consumers with greater pricing flexibility, or when content providers own some “superstar” content valued by every customer.

5.3.1 Price discrimination

Given that the platform is able to provide type-specific recommendations to its customers, it may also have the ability to price-discriminate according to these consumer types. When this is the case, the equilibrium in our benchmark scenario (where the platform does not use its recommendation system strategically) remains unchanged. Indeed, given that participating customers are homogeneous in their valuation v when facing their respective optimal content-mix, the platform is unable to increase its profit further through price discrimination.

However, the platform may use its recommendation system strategically and also adapt its pricing strategy as a function of customers’ types, setting a price $P(x)$ to participating customers of type x . Consumers of type x would participate, as long as $v - P(x) - t[\rho(x) - x]^2 \geq 0$. Therefore, the platform would set $P(x) = v - t[\rho(x) - x]^2$ for a given recommendation system, thus leaving consumers of type x with zero surplus.

In this case, the margin earned by the platform over a participating customer of type x is given by $P(x) - r_B + \Delta\rho(x)$. An increase in $\rho(x)$ would lower the platform’s costs for this customer type, but would also limit the price $P(x)$ it could charge. The optimal recommendation determined by the platform is thus $\rho(x) = x + \Delta/(2t)$, $\forall x \leq 1 - \Delta/(2t)$, and $\rho(x) = 1$ otherwise. The resulting equilibrium price is given by $P^*(x) = v - \Delta^2/(4t)$, $\forall x \leq 1 - \Delta/(2t)$, and $P^*(x) = v - t(1 - x)^2$ otherwise. The impact of the ability to price-discriminate is summarized in the following Proposition.

Proposition 7. *When the platform can price-discriminate according to consumers’ types, the recommendation bias is (weakly) larger. Moreover, the platform earns greater profit and consumers are left with zero surplus.*

Proof. See Appendix. □

Unsurprisingly, the ability for the platform to price-discriminate increases its profit, to the detriment of consumers. More interesting, however, is the fact that the platform sets a recommendation bias which is larger under price discrimination, therefore reducing the cost of serving each customer type. The platform also adjusts its pricing, and prices are lower under price discrimination than under uniform pricing for “low-type” customers (i.e., for customers with type $x < 2t/(\Delta + 2t)$, as determined by equation (5)). By contrast, customers which could extract a positive surplus under uniform pricing face higher prices under price discrimination.

5.3.2 Superstar content

In the baseline model, we have assumed that no content provider provides content that is essential for all consumers. In reality, content providers may own “superstar” content that all consumers would like to include in their content-mix. In this extension, we thus consider the case where some superstar content is owned by the content providers.

We assume that consumers’ types x are uniformly distributed over the interval $[\bar{x}, 1 - \bar{x}]$, with $\bar{x} \in (0, 1/2)$, and denote by $f(x) = 1/(1 - 2\bar{x})$ the density function. A consumer of type $x \in [\bar{x}, 1 - \bar{x}]$ has an optimal mix consisting of $x\%$ of content A and $(1 - x)\%$ of content B . Hence, even the consumers who like content A the least (i.e., those with type $x = \bar{x}$) include some content from A (i.e., its “superstar” content) in their optimal mix. Similarly, the consumers who like content B the least (i.e., those with type $x = 1 - \bar{x}$) include the superstar content from B in their mix. Below, we show that in this variation of our baseline setting, the platform may be able to capture the entire surplus created with strategic recommendations, and consumer welfare may not increase.

As in the main model, for any consumer type $x \in [\bar{x}, 1 - \bar{x}]$, the profit-maximizing recommendation system is given by $\rho(x) = \min\{1, x + \sqrt{(v - P)/t}\}$. We then define the threshold $\dot{x}(P) \equiv 1 - \sqrt{(v - P)/t}$ such that $\rho(x) = 1$ for $x \geq \dot{x}$, and $\rho(x) = x + \sqrt{(v - P)/t}$ for $x < \dot{x}$. Two cases can arise, depending on whether $\dot{x} \leq 1 - \bar{x}$ or the reverse. Solving for the equilibrium in each of these two cases, we obtain the following result.²⁷

Proposition 8. *When content providers offer superstar content and $\Delta > 0$, recommendations are biased in equilibrium. In addition:*

- *If $\bar{x} \leq \Delta/(2t)$, the equilibrium subscription price is $P^* = v - t\Delta^2(1 - \bar{x})^2/[\Delta + 2t(1 - 2\bar{x})]^2$, and recommendations are given by $\rho(x) = \min\{1, x + \Delta(1 - \bar{x})/[\Delta + 2t(1 - 2\bar{x})]\}$, which leaves some positive surplus for the participating customers with type $x \in [\dot{x}, 1 - \bar{x}]$.*
- *If $\bar{x} > \Delta/(2t)$, the equilibrium subscription price is $P^* = v - \Delta^2/(4t)$, recommendations are given by $\rho(x) = x + \Delta/(2t)$, and all consumers are left with zero surplus.*

Proof. See Appendix. □

This proposition shows that when consumers view some content as essential in their content-mix – for instance, when the providers offer “superstar” content that is valued by all the potential users of the platform – the platform may bias its recommendations and price its service such that all consumers are left with zero surplus in equilibrium. This is the case when providers A and B offer a large share of superstar content (i.e., \bar{x} is high) and/or

²⁷We verified that the platform has no incentive to exclude some consumers or to set a full bias. The proof is available in the web appendix.

when the difference between royalty rates is small. In this case, the platform captures the entire surplus created through the use of strategic recommendations (i.e., compared to the benchmark scenario). This implies that, for a given average royalty rate, the platform can earn greater profits when content providers own some “superstar” content.

6 Managerial Implications

Our analysis bears several important managerial implications. First, at given royalties, a platform may gain by steering consumers towards cheaper content, even though consumer valuation for the service may decrease as a result. This relates to previous work which aimed at putting the emphasis not only on customer retention but, more generally, on profitability (see, e.g., Bodapati (2008)). In our model, providing users with their preferred content-mix may not be the optimal strategy for the platform.²⁸ Recommendation systems need not always satisfy a purely customer-centric view, and can be tailored to take into account other drivers of profitability, such as content costs. For instance, Netflix reduced its costs by steering recommendations towards older DVDs it had already in stock, as explained by Shih and Kaufman (2014). Our analysis highlights the trade-off for the platform between revenue-maximization, which requires to provide optimal recommendations to users, and cost-minimization, which is achieved by offering biased personalized recommendations. Our analysis suggests that platform managers should put the balance towards cost-minimization when content costs are highly heterogeneous and/or consumers do not have strong preferences for their content-mix.

Second, streaming platforms which use their recommendation systems strategically can leverage their ability to steer consumption in order to reduce the market power of content providers. As mentioned above, the online-radio Pandora and the indie-label coalition Merlin engaged into a contract in which Merlin accepted reduced rates in exchange of an increased performance – i.e., a greater recommendation rate of its titles. This provides a good illustration of the impact of a strategic recommendation system on the business-to-business arrangements between the platform and content providers. When the platform can negotiate similar contracts with all content providers, or when these providers anticipate that the platform will make a strategic use of its recommendation system, important savings can be realized by the platform. This is because a strategic recommendation system can enable competition between content providers.

Third, a streaming platform can decide to integrate into the content production segment, not only to access content at a cheaper price but also to credibly commit to steering

²⁸We have also demonstrated that steering can be an optimal strategy even when users are able to search for their favorite content.

towards its own in-house productions. In our model, vertical integration (i) coordinates the supply-chain with respect to integrated content, but (ii) weakly increases the double-marginalization problem for content sourced from the remaining, independent content provider. Indeed, the independent content provider’s best response to integration is to set high royalty rates while facing a low demand, because of the recommendation bias which arises in equilibrium. As the independent content provider’s profit decreases with the platform’s integration, further strategic effects, such as foreclosure, could occur. In this respect, it is interesting to note that, in recent years, Netflix has invested massively in its own content and has announced a goal of filling 50% of its catalog with in-house productions.²⁹ In the music industry, Spotify is also moving towards more integration with content producers.³⁰ Given the ability of these platforms to steer consumption through strategic recommendations, our model predicts that this recent trend towards more integration into content production could have an important impact on contract (re)negotiations with independent content providers, to the ultimate benefit of the platforms.

7 Conclusion

In this paper, we have analyzed how a streaming platform which is constrained to set a uniform subscription price to all participating consumers and no usage fee can bias its personalized recommendations in order to steer consumption towards its most profitable content. When the royalty rates paid by the platform to content providers are set exogenously, customers face biased recommendations: each consumer type is recommended a larger percentage of the content which is cheaper for the platform than in its optimal content-mix. Consumers obtain personalized recommendations, but generally face the same bias (i.e., the same difference between the recommendation and their optimal content-mix) in equilibrium.

When content providers are able to set their royalty rates in the first period of the game, we showed that recommendations can lead to greater profits for the platform when used strategically, that is, when the recommendation system is optimized considering the platform’s overall profitability. Through the ability to steer consumption, the (threat of a biased) recommendation system can reduce royalty rates to a level lower than that obtained when the recommendation system only aims at maximizing consumers’ utility.

Our results are robust to the case where participating customers can search themselves for their favorite content on the platform, and also extend to the scenario where the platform produces its own content. Finally, we also investigated the cases where the platform can

²⁹Source: McCormick (2016).

³⁰See Nicolaou (2018).

price-discriminate according to consumers' types, and where content providers own some "superstar" content which all consumers want in their optimal content-mix.

References

- Adams, William, and Janet L. Yellen.** 1976. "Commodity Bundling and the Burden of Monopoly." *Quarterly Journal of Economics*, 90(3): 475–498.
- Aguiar, Luis, and Joel Waldfogel.** 2018. "Platforms, Promotion, and Product Discovery: Evidence from Spotify Playlists." JRC Digital Economy Working Paper 2018-04.
- Anderson, Simon P., and Damien J. Neven.** 1989. "Market efficiency with combinable products." *European Economic Review*, 33(4): 707–719.
- Ansari, Asim, Skander Essegai, and Rajeev Kohli.** 2000. "Internet Recommendation Systems." *Journal of Marketing Research*, 37(3): 363–375.
- Armstrong, Mark, and Jidong Zhou.** 2011. "Paying for Prominence." *The Economic Journal*, 121(556): F368–F395.
- Bakos, J. Yannis.** 1997. "Reducing Buyer Search Costs: Implications for Electronic Marketplaces." *Management Science*, 43(12): 1676–1692.
- Bardey, David, Denis Gromb, David Martimort, and Jérôme Pouyet.** 2018. "Controlling Sellers Who Provide Advice: Regulation and Competition." Mimeo.
- Belleflamme, Paul.** 2016. "The Economics of Digital Goods: A Progress Report." *Review of Economic Research on Copyright Issues*, 13(2): 1–24.
- Bodapati, Anand V.** 2008. "Recommendation Systems with Purchase Data." *Journal of Marketing Research*, 45(1): 77–93.
- Brynjolfsson, Erik, Yu (Jeffrey) Hu, and Duncan Simester.** 2011. "Goodbye Pareto Principle, Hello Long Tail: The Effect of Search Costs on the Concentration of Product Sales." *Management Science*, 57(8): 1373–1386.
- Burguet, Roberto, Ramon Caminal, and Matthew Ellman.** 2015. "In Google we trust?" *International Journal of Industrial Organization*, 39: 44–55.
- Calvano, Emilio, and Bruno Jullien.** 2018. "Can we trust the algorithms that recommend products online? A theory of biased advice with no pecuniary incentives and lab evidence." Mimeo.
- Carr, David.** 2013. "Giving Viewers What They Want." *New York Times*, February 24, www.nytimes.com/2013/02/25/business/media/for-house-of-cards-using-big-data-to-guarantee-its-popularity.html.
- Casadesus-Masanell, Ramon, and Hanna Halaburda.** 2014. "When Does a Platform Create Value by Limiting Choice?" *Journal of Economics & Management Strategy*, 23(2): 259–293.

- Chung, Tuck Siong, Roland T. Rust, and Michel Wedel.** 2009. "My Mobile Music: An Adaptive Personalization System for Digital Audio Players." *Marketing Science*, 28(1): 52–68.
- Datta, Hannes, George Knox, and Bart J. Bronnenberg.** 2018. "Changing Their Tune: How Consumers' Adoption of Online Streaming Affects Music Consumption and Discovery." *Marketing Science*, 37(1): 5–21.
- de Cornière, Alexandre, and Greg Taylor.** 2014. "Integration and search engine bias." *RAND Journal of Economics*, 45(3): 576–597.
- de Cornière, Alexandre, and Greg Taylor.** 2016. "A Model of Biased Intermediation." CEPR Discussion Paper 11457.
- Drugov, Mikhail, and Doh-Shin Jeon.** 2017. "Vertical Integration and Algorithm Bias." Mimeo.
- Fleder, Daniel, and Kartik Hosanagar.** 2009. "Blockbuster Culture's Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity." *Management Science*, 55(5): 697–712.
- Gabszewicz, Jean J., Didier Laussel, and Nathalie Sonnac.** 2004. "Programming and Advertising Competition in the Broadcasting Industry." *Journal of Economics & Management Strategy*, 13(4): 657–669.
- Hagiu, Andrei, and Bruno Jullien.** 2011. "Why do intermediaries divert search?" *RAND Journal of Economics*, 42(2): 337–362.
- Hagiu, Andrei, and Bruno Jullien.** 2014. "Search diversion and platform competition." *International Journal of Industrial Organization*, 33: 48–60.
- Hunold, Matthias, and Johannes Muthers.** 2017. "Resale price maintenance and manufacturer competition for retail services." *RAND Journal of Economics*, 48(1): 3–23.
- Inderst, Roman, and Marco Ottaviani.** 2012. "Competition through Commissions and Kickbacks." *American Economic Review*, 102(2): 780–809.
- Jeon, Doh-Shin, and Domenico Menicucci.** 2012. "Bundling and Competition for Slots." *American Economic Review*, 102(5): 1957–1985.
- Mangalindan, JP.** 2012. "Amazon's recommendation secret." *Fortune*, July 30, www.fortune.com/2012/07/30/amazons-recommendation-secret/.
- McAfee, R. Preston, John McMillan, and Michael D. Whinston.** 1989. "Multiproduct Monopoly, Commodity Bundling, and Correlation of Values." *Quarterly Journal of Economics*, 104(2): 371–383.
- McCormick, Rich.** 2016. "Netflix planning to fill half its catalog with originals in the next few years." *The Verge*, September 21, www.theverge.com/2016/9/21/12997058/netflix-originals-half-catalog-streaming.

- Nicolaou, Anna.** 2018. "Spotify shakes record labels by dealing directly with artists." *Financial Times*, June 15, www.ft.com/content/f1b27624-6e68-11e8-92d3-6c13e5c92914.
- Oestreicher-Singer, Gal, and Arun Sundararajan.** 2012. "Recommendation Networks and the Long Tail of Electronic Commerce." *MIS Quarterly*, 36(1): 65–83.
- Pathak, Bhavik, Robert Garfinkel, Ram D. Gopal, Rajkumar Venkatesan, and Fang Yin.** 2010. "Empirical Analysis of the Impact of Recommender Systems on Sales." *Journal of Management Information Systems*, 27(2): 159–188.
- Popper, Ben.** 2016. "Spotify's Discover Weekly reaches 40 million users and 5 billion tracks streamed." *The Verge*, May 25, www.theverge.com/2016/5/25/11765472/spotify-discover-weekly-40-million-users-5-billion-streams.
- Raskovich, Alexander.** 2007. "Retail buyer power through steering." *Economics Letters*, 96(2): 221–225.
- Ricci, Francesco, Lior Rokach, Bracha Shapira, and Paul B. Kantor,** ed. 2011. *Recommender Systems Handbook*. Springer.
- Shih, Willy, and Stephen Kaufman.** 2014. "Netflix." Harvard Business School Case 615-007.
- Sydell, Laura.** 2014. "Pandora's New Deal: Different Pay, Different Play." *npr.org*, November 26, www.npr.org/2014/11/26/366339553/pandoras-new-deal-different-pay-different-play.
- Thomes, Tim Paul.** 2013. "An economic analysis of online streaming music services." *Information Economics and Policy*, 25(2): 81–91.
- Vanderbilt, Tom.** 2013. "The Science Behind the Netflix Algorithms That Decide What You'll Watch Next." *Wired*, July 8, www.wired.com/2013/08/qq_netflix-algorithm/.
- Xiao, Bo, and Izak Benbasat.** 2007. "E-commerce Product Recommendation Agents: Use, Characteristics, and Impact." *MIS Quarterly*, 31(1): 137–209.

Appendices

Proof of Proposition 1. In the main text, we have characterized the equilibrium outcome when the market is covered and both contents, A and B , are recommended to consumers. The other cases which could arise are (i) a covered market with only the cheaper content recommended to consumers, (ii) a market where some consumers are excluded, with only the cheaper content recommended to consumers, and (iii) a market where some consumers are excluded and both contents A and B are recommended to consumers. We demonstrate below that these alternative strategies are dominated for the platform by the main strategy outlined in Proposition 1, where the platform covers the market and recommends both A and B .

First, we demonstrate that the platform would not cover the market and recommend only content A (the cheaper content) to participating consumers, that is, set $\rho(x) = 1$ for all $x \in [0, 1]$. Doing so would allow the platform to charge a maximum subscription price equal to $v - t$, which represents the gross utility of consumers of type $x = 0$, who dislike content A the most. Hence, the platform would earn a total profit of $v - t - r_A$, which is strictly lower than the profit obtained under a covered market where both A and B are provided to consumers. Therefore, it is a dominated strategy for the platform to sell to all consumers, while recommending content A to all of them.

Second, we demonstrate that it is a dominated strategy for the platform to sell to consumers with types in the interval $[\widehat{x}, 1]$, with $\widehat{x} \in (0, 1)$, while recommending $\rho(x) = 1$ (i.e., content A) to all of them. In this scenario, the platform's profit is given by $(1 - \widehat{x})(P - r_A)$, and the maximum subscription price the platform can set is constrained by the participation constraint of consumers of type $x = \widehat{x}$; i.e., it must be that $P \leq v - t(1 - \widehat{x})^2$. The platform hence sets the subscription fee $P = v - t(1 - \widehat{x})^2$ and then designs its recommendation system such that $\widehat{x} = 1 - \sqrt{(v - r_A)/(3t)}$, with $t > (v - r_A)/3$ for some consumers to be excluded by the platform. The resulting profit is always lower than the profit obtained under a covered market where both A and B are provided to consumers. Hence, it is a dominated strategy for the platform to sell to consumers with types in the interval $[\widehat{x}, 1]$, while recommending $\rho(x) = 1$ to all of them.

Third, and finally, we compute the equilibrium subscription price when the market is not covered and both products A and B are recommended to consumers. The platform's profit is then given by:

$$\Pi^U = \int_{\widehat{x}(P)}^1 (P - r_A) dx - \Delta \int_{\widehat{x}(P)}^{\widehat{x}(P)} (1 - \rho(x)) dx, \quad (10)$$

which develops as:

$$\Pi^U = \frac{1}{2\Delta} \left\{ P^2 - 2r_A P - r_A^2 + 2\Delta P \sqrt{\frac{v - P}{t}} + 2r_A \left[r_B - \Delta \left(1 + \sqrt{\frac{v - P}{t}} \right) \right] \right\}. \quad (11)$$

We show below that the platform never chooses to exclude consumers in equilibrium. We find that $\widehat{x}''(P) > 0$, that $\lim_{P \rightarrow -\infty} \widehat{x}'(P) < 0$ and that $\lim_{P \rightarrow v^-} \widehat{x}'(P) > 0$. Therefore, $\widehat{x}(P)$

is first decreasing and then increasing for $P \in [0, v)$. Furthermore, we find that $\widehat{x}(v) < 0$. Therefore, either $\widehat{x}(P)$ is always negative for $P \in [0, v)$ or it is first positive then negative. It follows that if $\widehat{x}(0) > 0$ (which occurs if and only if $r_B/\Delta > \sqrt{v/t}$), there is a unique $\widehat{P} > 0$ that is the solution of $\widehat{x}(\widehat{P}) = 0$, such that $\widehat{x}(P) > 0$ if and only if $P \leq \widehat{P}$. We find that:

$$\widehat{P} = r_B - \frac{\Delta^2 + \Delta \sqrt{\Delta^2 + 4t(v - r_B)}}{2t}. \quad (12)$$

We then know that the platform's profit function is defined in three parts: (i) if $P \leq r_A$, the platform makes losses if it sells subscriptions to consumers, so it shuts down and makes zero profit; (ii) if $P \in [r_A, \widehat{P}]$, the platform excludes some consumers and its profit is given by (10); (iii) if $P > \widehat{P}$, the platform covers the full market and its profit is given by (3).

We are now going to prove that the platform's profit is increasing over the range where it excludes some consumers (i.e., for $P \in [r_A, \widehat{P}]$).

Let $\Pi^{U(n)}(P)$ denote the n -th derivative of $\Pi^U(P)$ with respect to P . We find that $\Pi^{U(4)} < 0$, which shows that $\Pi^{U(3)}(P)$ is decreasing. Since $\Pi^{U(3)}(r_A) < 0$, then $\Pi^{U(3)}(P) < 0$ for all $P \in [r_A, \widehat{P}]$. This shows that $\Pi^{U(2)}(P)$ is decreasing. We then find that $\Pi^{U(2)}(r_A)$ can be either positive or negative and that $\lim_{P \rightarrow v^-} \Pi^{U(2)}(P) = -\infty$. Therefore, $\Pi^{U(1)}(P)$ is either (i) always decreasing, or (ii) first increasing, then decreasing. Since $\Pi^{U(1)}(r_A) > 0$ and $\lim_{P \rightarrow v^-} \Pi^{U(1)}(P) = -\infty$, in both cases (i) and (ii), it means that $\Pi^{U(1)}(P)$ is positive up to a value of P , and then negative. If we can then prove that $\Pi^{U(1)}(\widehat{P}) > 0$, it would show that $\Pi^{U(1)}(P) > 0$ for all $P \in [r_A, \widehat{P}]$, that is, $\Pi^U(P)$ is increasing over the relevant range.

We now demonstrate that the platform's profit Π^U is increasing at $P = \widehat{P}$, which will prove that it is increasing over the range where the platform finds it profitable to exclude some consumers.

We first show that the derivative of Π^U at $P = \widehat{P}^-$ is the same as the derivative of Π at $P = \widehat{P}^+$. To begin with, we find that:

$$\frac{\partial \Pi}{\partial P} - \frac{\partial \Pi^U}{\partial P} = \widehat{x}(P) (1 + \Delta b'(P)) + \widehat{x}'(P) [P - r_B + \Delta (b(P) + \widehat{x}(P))], \quad (13)$$

where $b(P) \equiv \sqrt{(v - P)/t}$. Since $\widehat{x}(\widehat{P}) = 0$, we have :

$$\left. \frac{\partial \Pi}{\partial P} - \frac{\partial \Pi^U}{\partial P} \right|_{P=\widehat{P}} = \widehat{x}'(\widehat{P}) [\widehat{P} - r_B + \Delta b(\widehat{P})] = -\Delta \widehat{x}'(\widehat{P}) \widehat{x}(\widehat{P}) = 0. \quad (14)$$

Therefore, the derivative of Π^U at $P = \widehat{P}^-$ is the same as the derivative of Π at $P = \widehat{P}^+$.

Second, the platform's profit under coverage is concave as:

$$\frac{\partial^2 \Pi}{\partial P^2} = \frac{-\Delta}{4t^2 b(P)^{3/2}} \leq 0, \quad (15)$$

and it is strictly concave if $\Delta > 0$. Furthermore, we find that $\widehat{P} < P^*$. Indeed, for a given Δ ,

\widehat{P} is increasing in r_B and:

$$\widehat{P}\Big|_{r_B=v} = v - \frac{\Delta^2}{t} < v - \frac{t\Delta^2}{(\Delta + 2t)^2} = P^*. \quad (16)$$

Since Π is concave and $\widehat{P} < P^*$, then \widehat{P} lies on the increasing part of Π , which implies that if $\Delta > 0$ then:

$$\frac{\partial \Pi^U}{\partial P}\Big|_{P=\widehat{P}} = \frac{\partial \Pi}{\partial P}\Big|_{P=\widehat{P}} > 0. \quad (17)$$

We have thus shown that the platform's profit is increasing in P over the range where the platform finds it profitable to exclude some consumers. Since the profit function is continuous, this means that P^* is the global maximum of Π . \square

Proof of Proposition 2. We first show that $r_A^* = r_B^* = t$ is an equilibrium when $v \geq t$. We see that (i) $\partial \pi_B / \partial r_B = -2t^2(r_A + r_B - 2t) / (\Delta + 2t)^3$ and (ii) $\partial \pi_A / \partial r_A = 1 - 4r_A t^2 / (\Delta + 2t)^3 - 2t^2 / (\Delta + 2t)^2$ both equal zero when $r_A = r_B = t$. Any unilateral deviation is unprofitable because $\partial \pi_B / \partial r_B \leq 0$, $\forall r_B \geq t$ when $r_A = t$, and $\partial \pi_A / \partial r_A \geq 0$, $\forall r_A \leq t$ when $r_B = t$. Besides, (iii) we always have $\partial^2 \pi_A / \partial r_A^2 < 0$, and we have (iv) $\partial^2 \pi_B / \partial r_B^2 < 0$ when $t > (2r_A + r_B)/4$, which is satisfied at $r_A = r_B = t$.

Then, we demonstrate that $r_A^* = r_B^* = v$ is an equilibrium when $t > v$. Indeed, when $t > v$, $r_A = r_B = t$ cannot be an equilibrium because the platform cannot charge more than v to its subscribers, and hence it would not cover its royalty costs and shut down. We then see that (i) $\partial \pi_B / \partial r_B$ and (ii) $\partial \pi_A / \partial r_A$ are both strictly positive at $r_A = r_B = v$. Hence, both firms A and B would like to increase (unilaterally) their respective royalty rates (but they cannot under Assumption 1). Besides, (iii) we always have $\partial^2 \pi_A / \partial r_A^2 < 0$, and we have (iv) $\partial^2 \pi_B / \partial r_B^2 < 0$ when $t > (2r_A + r_B)/4$, which is satisfied at $r_A = r_B = v$.

Finally, we show that there can be no other equilibrium. First, when $t > v$, we see that (i) $\partial \pi_B / \partial r_B$ and (ii) $\partial \pi_A / \partial r_A$ are both positive at any $r_A \leq r_B < v < t$. Second, when $v \geq t$ and $r_A > t$, we observe that (i) $\partial \pi_B / \partial r_B < 0$. Moreover, (ii) $\partial \pi_A / \partial r_A < 0$ at $r_B = r_A$. Third, when $v \geq t$, $r_A \leq t$ and $v > 2t - r_A$, we see that (i) $\partial \pi_B / \partial r_B > 0 \Leftrightarrow r_B < 2t - r_A$. Moreover, (ii) $\partial \pi_A / \partial r_A > 0$ at $r_B = 2t - r_A > t$. Fourth, and finally, when $v \geq t$, $r_A \leq t$ and $v < 2t - r_A$, we can see that (i) $\partial \pi_B / \partial r_B > 0$ for all $r_B \leq v$, and also that (ii) $\partial \pi_A / \partial r_A > 0$ at $r_B = v$. Hence, at least one firm always has a unilateral incentive to deviate, except at the two equilibria detailed above. \square

Proof of Proposition 4. At the first stage of the game, provider A 's and provider B 's profits are

$$\pi_A = \begin{cases} r_A \left[1 - \frac{2t^2}{(\Delta + 2t)^2} \right] & \text{if } s \geq t\Delta^2 / (\Delta + 2t)^2 \\ \frac{r_A}{2} \left(1 - \frac{s}{t} + 2\sqrt{\frac{s}{t}} \right) & \text{if } s < t\Delta^2 / (\Delta + 2t)^2 \end{cases} \quad (18)$$

and

$$\pi_B = \begin{cases} r_B \left[\frac{2t^2}{(\Delta + 2t)^2} \right] & \text{if } s \geq t\Delta^2 / (\Delta + 2t)^2 \\ \frac{r_B}{2} \left(1 - \sqrt{\frac{s}{t}} \right)^2 & \text{if } s < t\Delta^2 / (\Delta + 2t)^2 \end{cases} \quad (19)$$

respectively. Note that the condition $s \geq t\Delta^2 / (\Delta + 2t)^2$ is satisfied if Δ is low enough.

We start by checking whether the equilibrium given by Proposition 2 still holds when consumers can search for content. If $v < t$, the equilibrium is such that $r_A^* = r_B^* = v$. Content providers already capture all surplus in equilibrium and cannot increase their royalty rates.

If $v \geq t$, the equilibrium in Proposition 2 is such that $r_A^* = r_B^* = t$. From the proof of Proposition 2, we know that the profit of content provider B is decreasing for a small increase in r_B above t , if $r_A = t$. However, from the expression of B 's profit in (19), we see that B 's profit becomes increasing if Δ , and hence r_B , is large enough. Therefore, there is the possibility that provider B could increase its profit with a large deviation, by changing its royalty rate r_B from t up to v . This deviation can be profitable only if B reaches the second branch of its profit function, that is, if $s < t[(1 - t/v)/(1 + t/v)]^2 \equiv \bar{s}$. If this is the case, we find that B would obtain a profit of $r_B \ddot{x} [1 - \rho(0)] / 2 = (v/2)(1 - \sqrt{s/t})^2$. This deviation is profitable for B if and only if $(v/2)(1 - \sqrt{s/t})^2 > t/2$, which is equivalent to $s < t(1 - \sqrt{t/v})^2 \equiv \hat{s} \leq \bar{s}$.

Therefore, if $v \geq t$ and $s \geq \hat{s}$, there is no profitable deviation for B and the equilibrium royalty rates are still given by Proposition 2, i.e., we have $r_A^* = r_B^* = t$.

If $v \geq t$ and $s < \hat{s}$, content provider B has an incentive to set $r_B = v$, if $r_A = t$. In turn, if $r_B = v$, content provider A clearly has an incentive to raise its royalty rate, up to $\bar{r}_A \equiv v - 2t\sqrt{s}/(\sqrt{t} - \sqrt{s}) \in [t, v]$, at which point we have $s = t\Delta^2/(\Delta + 2t)^2$, as the platform will not modify its recommendation system for any r_A below this threshold. Moreover, A may even have an incentive to raise its royalty rate above \bar{r}_A if this maximizes its profit when $r_B = v$. At this rate set by A , content provider B , in turn, would have an incentive to decrease its royalty rate as $\partial\pi_B/\partial r_B < 0$, for all r_B such that $s \geq t\Delta^2/(\Delta + 2t)^2$. Then, following the proof of Proposition 2 for the case where $v \geq t$ and $r_A > t$, we see that both firms would have incentives to lower their royalty rates down to t . This would trigger again a large deviation by content provider B up to $r_B = v$, which shows that there cannot be any Nash equilibrium in pure strategies when $v \geq t$ and $s < \hat{s}$. \square

Proof of Proposition 5. When $v \geq 2t$, the content provider B has an incentive to set $r_B = 2t$. In this case, using $r_A = 0$ and $r_B = 2t$ together with Proposition 1 proves the results. Similarly, when $v < 2t$, content provider B is limited by Assumption 1 and sets $r_B = v$. Again, using $r_A = 0$ and Proposition 1 proves the results. \square

Proof of Proposition 7. We show that the recommendation bias is (weakly) larger under price discrimination than under uniform pricing. First, consumers with types $x \leq 1 - \Delta/(2t)$ face a mix of content from A and B both under uniform pricing and price discrimination. Clearly, the bias they face under price discrimination, $\Delta/(2t)$, is greater than under uniform pricing, $\Delta/(\Delta + 2t)$. Second, consumers with type $1 - \Delta/(2t) < x \leq 2t/(\Delta + 2t)$ face a recommendation of $\rho(x) = 1$ under price discrimination, and $\rho(x) = x + \Delta/(\Delta + 2t)$ under uniform pricing. Again, the bias under discrimination, $1 - x$, is greater than under uniform pricing, $\Delta/(\Delta + 2t)$. Finally, consumers of type $x > 2t/(\Delta + 2t)$ face the same bias under price discrimination and uniform pricing, as $\rho(x) = 1$ under both regimes.

In addition, for a given pair of royalty rates, the platform is able to extract the entire surplus under price discrimination, thus simultaneously increasing its profit and reducing consumer surplus to zero, compared to the case with uniform pricing. \square

Proof of Proposition 8. We first investigate the case where $\dot{x} \in (\bar{x}, 1 - \bar{x})$. In this case, the

analysis is very similar to that of the baseline model. The platform's profit is given by:

$$\Pi = \int_{\bar{x}}^{1-\bar{x}} (P - r_A) f(x) dx - \Delta \int_{\bar{x}}^{\dot{x}(P)} [1 - \rho(x)] f(x) dx, \quad (20)$$

with $f(x) = 1/(1 - 2\bar{x})$. In equilibrium, we obtain the subscription price

$$P^* = v - \frac{t\Delta^2(1 - \bar{x})^2}{[\Delta + 2t(1 - 2\bar{x})]^2},$$

which gives

$$\rho(x) = \min \left\{ 1, x + \frac{\Delta(1 - \bar{x})}{\Delta + 2t(1 - 2\bar{x})} \right\} \text{ and } \dot{x} = 1 - \frac{\Delta(1 - \bar{x})}{\Delta + 2t(1 - 2\bar{x})}.$$

As in the baseline model, strategic recommendations lead to positive surplus for participating customers with type $x \in [\dot{x}, 1 - \bar{x}]$. This result holds if and only if $\dot{x}(P^*) \in (\bar{x}, 1 - \bar{x})$. We find that

$$\dot{x} = \bar{x} + \frac{2t(1 - \bar{x})(1 - 2\bar{x})}{\Delta + 2t(1 - 2\bar{x})} > \bar{x},$$

and that

$$(1 - \bar{x}) - \dot{x} = \frac{(\Delta - 2t\bar{x})(1 - 2\bar{x})}{\Delta + 2t(1 - 2\bar{x})}$$

has the sign of $\Delta - 2t\bar{x}$. Therefore, we have $\dot{x} \in [\bar{x}, 1 - \bar{x}]$ if and only if $\bar{x} < \Delta/(2t)$.

We now consider the case where $\dot{x} > 1 - \bar{x}$. The platform sets its subscription price P in order to maximize its profit, which is given by:

$$\Pi = \int_{\bar{x}}^{1-\bar{x}} (P - r_A) f(x) dx - \Delta \int_{\bar{x}}^{1-\bar{x}} [1 - \rho(x)] f(x) dx. \quad (21)$$

The main difference with the previous case lies in the fact that in the second integral, $1 - \bar{x}$, which is exogenous and does not depend on the subscription price, replaces $\dot{x}(P)$.

At the equilibrium, we obtain that $P^* = v - \Delta^2/(4t)$, which gives $\rho(x) = x + \Delta/(2t)$. We find that $\dot{x}(P^*) = 1 - \Delta/(2t)$, and therefore these results hold if $\dot{x}(P^*) > 1 - \bar{x}$, that is, if $\bar{x} > \Delta/(2t)$. In this case, all participating consumers are left with zero surplus. Therefore, if $\bar{x} > \Delta/(2t)$, a strategic recommendation system does not generate any positive consumer surplus in equilibrium. \square