WORKING PAPERS

N° 1742

April 2026

“Dynamic Recommendation Bias”

Mikhail Drugov and Doh-Shin Jeon

Toulouse
School of
Economics

Dynamic Recommendation Bias[*]

Mikhail Drugov[†] Doh-Shin Jeon[‡]

This version: April 29, 2026

Abstract

This paper studies the incentives of a subscription-funded platform that offers both proprietary and third-party content to bias its recommendations about which content users should consume. Consistent with Netflix's practice, we consider fixed-fee bargaining between the platform and a content provider, which eliminates any static incentive to bias recommendations. However, our dynamic model identifies two distinct incentives to bias recommendations: improving the platform's future bargaining position and increasing users' expected surplus. The former favors first-party content, while the latter favors the *ex ante* superior content. As a result, biased recommendations may lead to either self-preferencing or third-party preferencing.

Keywords: Recommendation, Platform, Algorithm, Signal Jamming.

JEL codes: D83, L42.

---

[*] We thank Michael Grubb, Bruno Jullien, Margaret Meyer, Patrick Rey and Sandro Shelegia and participants of various seminars and conferences for helpful comments. This paper was previously circulated under the title "Vertical Integration and Algorithm Bias". Drugov gratefully acknowledges grant PID2024-158416NB-I00 funded by Ministerio de Ciencia, Innovación y Universidades. Jeon acknowledges funding from ANR under grant ANR-17-EURE-0010 (Investissements d'Avenir program).

[†] Universitat Autònoma de Barcelona and Barcelona School of Economics (BSE), New Economic School, and CEPR, `mikhail.drugov@uab.cat`.

[‡] Toulouse School of Economics, University of Toulouse Capitole, `dohshin.jeon@tse-fr.eu`.

1 Introduction

> *Since it is very difficult even for experts to evaluate search engines, search engine bias is particularly insidious.*
>
> Sergey Brin and Larry Page

Recommendation algorithms are at the heart of many online platforms that deal with a large amount of information or content or a large number of products. More than 80% of streaming on Netflix is done following the platform's recommendations (Gomez-Uribe and Hunt, 2015). Among the followers of the top 1000 public playlists on Spotify, 90% follow those curated by Spotify (Aguiar, Waldfogel and Waldfogel, 2021).

Concerns about the potential bias of recommendation algorithms used by the platforms have been voiced repeatedly. Echoing the Google founders in the epigraph, the former European Competition Commissioner Vestager (2017), said "The trouble is, it's not easy to know exactly how those algorithms work. How they've decided what to show us, and what to hide." Some well-known examples of recommendation biases—typically taking the form of self-preferencing—include the Google shopping case (European Commission, 2017), Apple Appstore's biased search (Nicas and Collins, 2019), Amazon's bias towards its own products (European Commission, 2022; Farronato, Fradkin and MacKay, 2023). Spotify was discovered to bias its playlists in favor of independent-label music (Aguiar, Waldfogel and Waldfogel, 2021) and Netflix is known for actively promoting its original content.

This paper focuses on the platforms such as Netflix and Disney+, which use a subscription-funded business model and provide both first-party and third-party content. It addresses the following question: What is the mechanism potentially leading to their biased recommendations? In a static setting, if a platform's marginal costs of providing others' content are higher than those of providing its own content because of royalties, the platform will bias its recommendation towards its own content (Bourreau and Gaudin, 2022).[1] However, since the bias reduces the users' utility and thereby the subscription price that they are ready to pay, the platform and the content providers have a clear incentive to replace distortionary marginal pricing with a lump-sum licensing contract, so that the self-preferencing bias is corrected. Indeed, Netflix uses fixed fees to license third-party

---

[1] Even if the platform shows only third party content, it has an incentive to bias recommendations toward the content charging lower royalties.

content.[2]

Yet, in a dynamic setting the situation is quite different: Even with a lump-sum payment, the platform may have incentives to bias its recommendations. To identify the dynamic incentives, we build a two-period model with two main features which eliminate static incentives to bias recommendations: The platform makes a lump-sum payment to a third-party content provider and consumers pay subscription fees. In particular, in each period, before the platform chooses its recommendation bias, the platform and a single third-party content provider negotiate a lump-sum payment from the platform to host the provider's content during that period, while consumers pay subscription prices. Determining the payment requires the parties in the negotiation to estimate users' demand for the platform when it hosts both contents and the demand when it hosts its own content only—the latter serves as the platform's outside option in the negotiation. At the beginning of the first period, the users do not know which content is best for them; the platform, with its advanced algorithms and big data, knows it better and makes personalized recommendations. After the first-period lump-sum payment and the subscription price have been set, the platform has an incentive to bias its *first-period* recommendations to maximize its *second-period* profit. We identify two distinct incentives to bias recommendations: improving the platform's outside option and increasing users' expected surplus (from the second-period consumption). We call the former the outside-option effect and it generates biases toward first-party content; we call the latter the perceived-surplus effect and it generates biases toward the *ex ante* superior content.[3] The relative weight of the second effect is equal to the platform's bargaining power with respect to the content provider.

If the platform's content is of better quality, the two incentives are aligned and the equilibrium bias is towards its own content, leading to self-preferencing. If, instead, the other content is of higher quality, the equilibrium bias might be in favor of the other content, leading to third-party preferencing. This is more likely if the platform has a higher bargaining power. Indeed, in this case the platform is more concerned about increasing

---

[2] Netflix, Inc., Annual Report (Form 10-K) for the fiscal year ended Dec. 31, 2025, p. 28. `https://www.sec.gov/Archives/edgar/data/1065280/000106528026000034/nflx-20251231.htm` ("The content licenses are for a fixed fee and specific windows of availability.")

[3] An early reference for the outside-option effect is Klein, Crawford and Alchian (1978) and it has a prominent role in the property rights theory of the firm (Grossman and Hart, 1986; Hart and Moore, 1990). The perceived-surplus effect is sometimes called "pandering" and was studied in various settings. See, for example, Brandenburger and Polak (1996) where managers side with the prevailing market opinion, Heidhues and Lagerlöf (2003) for candidates' communication to voters during elections and Che, Dessein and Kartik (2013) for an abstract setting of an agent advising the principal. We are not aware of any model with this effect in the context of recommendation by a platform.

the (perceived) value of the platform hosting both contents than about investing in its outside option. A straightforward prediction is that as the platform's content becomes better, its recommendation is more likely to exhibit self-preferencing.

Related literature This paper belongs to the large literature on vertical foreclosure, see, e.g. Rey and Tirole (2007) for a review of earlier literature. A lot of the focus in the more recent literature has been on applying these ideas and developing new ones in relation to digital platforms, see, e.g., Motta (2023). In this recent literature, vertical foreclosure often takes the form of "self-preferencing", that is, the integrated platform biasing its search results and recommendations to favor its own services or products, or making the access to the platform more difficult for third parties (see Kittaka, Sato and Zennyo, 2023, for a survey of the literature). Early research (de Cornière and Taylor, 2014; Hagiu and Jullien, 2011, 2014) focused on "search bias", that is, bias by a search engine—with Google being the main example. Newer literature studies the incentives and the remedies for the "dual-mode" platforms—such as Amazon—which sell both their own products and those of third-party sellers, see de Cornière and Taylor (2019), Hagiu, Teh and Wright (2022), etc. As previously described, Bourreau and Gaudin (2022) study (subscription-funded) streaming platforms' incentive to bias recommendations when content providers charge different royalties. To the best of our knowledge, no paper has isolated platforms' dynamic incentives to bias recommendations. We contribute to the literature on platforms' recommendation bias or self-preferencing by identifying dynamic incentives to bias recommendations which can lead to either self-preferencing or third-party-preferencing. Although we focus on subscription-funded platforms to isolate the dynamic incentives, these incentives are also relevant for other platforms, since consumers often engage in repeated relationships with them.

The recommendation bias generated by what we call the perceived-surplus effect seems to be related to the popularity bias identified in diverse contexts by the experimental literature on recommendation systems, whereby already-popular items receive disproportionate exposure through recommendations, often at the expense of niche or less popular items. Fleder and Hosanagar (2009), Lee and Hosanagar (2019) and Calvano et al. (2026) find that the use of traditional recommendation systems such as collaborative filters is associated with a decrease in sales diversity. Abdollahpouri et al. (2021) show experimentally, using movie-recommendation data, that many recommendation algorithms produce lists that are extremely concentrated on popular items, even for users whose preferences include long-tail items. Calvano et al. (2026) explain this by the fact that—when the data quality is relatively poor—the algorithm overestimates the similarity between users and hence, focuses on the preferences of the median consumer. Celdir, Cho and Hwang (2024) develop a theoretical framework and provide empirical evidence of popularity bias in the context of online dating platforms. Our paper contributes to the literature on the popularity bias by providing a possible theoretical mechanism for the bias.

The rest of the paper is organized as follows. Section 2 introduces the model. The equilibrium is characterized in Section 3. Section 4 concludes. All proofs are contained in Appendix.

2 Model

There are two firms, firm A and firm B, each with its own content. Firm A is vertically integrated with a monopoly distribution network, i.e., a "platform". Firm B can distribute its content only through A's network. The distribution costs are zero for both contents.

There are two periods. There is a mass one of users who are located on $[0, 1]$. We assume that each user's location is perfectly correlated across the two periods and the location is uniformly distributed over $[0, 1]$, which is common knowledge. However, users do not know their realized locations. Content A is located at 0 and content B is located at 1.[4] Users incur linear transportation costs: A user located at x obtains utility $v_A - \tau x$ and $v_B - \tau(1-x)$ from content A and B, respectively, where $\tau > 0$ is a cost parameter and v_i is gross utility from content $i = A, B$.[5]

We assume that each period, conditional on consuming either content, a user has no demand for the other content. We also assume

Assumption 1 *(i)* $|v_A - v_B| < \tau$, *(ii)* $v_A > \tau$.

The first part of Assumption 1 guarantees a positive ex ante market share for each content. The second part of Assumption 1 is the standard full coverage assumption when only content A is available: Any user gets a positive gross surplus from consuming content A. The location of the marginal user $\hat{x}$ who is indifferent between consuming content A

---

[4] In fact, content A of the first period is different from content A of the second period. For simplicity, we call them content A. The same remark applies to content B. What really matters for our analysis is the *relative* location of users and the contents. We assume that the *relative* locations of users and the contents are perfectly correlated in the two periods. Qualitative results of our model hold if they are only imperfectly correlated.

[5] Our analysis does not require the same firm to own content B in both periods.

and content B is found from

$$v_A - \tau\hat{x} = v_B - \tau\left(1 - \hat{x}\right),$$

yielding

$$\hat{x} = \frac{1}{2} + \frac{\Delta}{2\tau}, \tag{1}$$

where $\Delta \equiv v_A - v_B$. By Assumption 1(i), $\hat{x} \in (0, 1)$.

We assume that each period $t = 1, 2$, firm A observes each user's location and recommends one of the two contents when both are available on A's distribution network. Hence, A's strategy is a mapping from the location of each user to a personalized recommendation $r_t(x) \in \{A, B\}$ with $t = 1, 2$. Firm A uses a cut-off strategy as follows

$$r_t(x) = \begin{cases} A & \text{if } x \leq \hat{x} + b_t \\ B & \text{if } x > \hat{x} + b_t \end{cases},$$

where b_t is the recommendation bias, which is not observed by the users.[6] In what follows, we call it the recommendation strategy with bias b_t. If $b_t > 0$, the recommendation is biased toward content A; if $b_t < 0$, the recommendation is biased toward content B.

In the beginning of each period, the two firms bargain over the fixed fee that firm A should pay to firm B in order to be able to distribute the latter's content. The stage game is the following. In each period $t = 1, 2$,

Stage 1 Firms A and B bargain over the fixed fee f_t that firm A pays to firm B for the right to distribute the latter's content in period t. If the bargaining fails, firm A distributes only its own content while firm B obtains zero. The outcome is given by the Nash bargaining solution, with $\alpha \in [0, 1]$ being A's bargaining power.

Stage 2 Firm A sets subscription price p_t for the users. The users decide whether to subscribe or not.

Stage 3 Upon subscription, users' locations are realized (in the first period as we assume that a user's second-period location is identical to her first-period one). Firm A observes each user's location and makes a personalized recommendation of content when both contents are available.

---

[6] As Brin and Page (1998) write, "For example, a search engine could add a small factor to search results from "friendly" companies, and subtract a factor from results from competitors". The cut-off strategy is the best for both the firms and the users since it is the least distortionary.

In the first period, as all users are ex ante identical, firm A charges the same subscription price to all of them. However, at the beginning of the second period, users are heterogeneous depending on the recommendation they received in the first period. In what follows, we assume that firm A can price discriminate users depending on the recommendation they previously received. This amounts to assuming full surplus extraction.

The solution concept is perfect Bayesian equilibrium (PBE). In a PBE, (i) the subscription prices of firm A are optimal given users' conjectures about recommendation biases of firm A and, in the second period, given the first-period recommendations; (ii) the subscription decisions of users are optimal given the subscription prices, their conjectures about recommendation biases of firm A and, in the second period, the first-period recommendations; (iii) the recommendation biases of firm A are optimal given users' conjectures about them; (iv) users' conjectures about recommendation biases are correct; (v) the Nash bargaining solution uses the correct recommendation biases. If there are multiple equilibria, we focus on the one that maximizes firm A's profit which is

$$\Pi_1^A + \delta \Pi_2^A,$$

where Π_t^A is firm A's profit in period $t = 1, 2$ and $\delta > 0$ is a common discount factor, which can be larger than one.

Let us now consider the benchmark model with only one period. There is an equilibrium without any bias and this equilibrium results in the highest profit for firm A. Indeed, consider stage 3 when firm A decides on the recommendation bias b. The users have already paid the subscription price at stage 2. Therefore, firm A is indifferent between recommending A and recommending B for any user and hence, it is indifferent with respect to the sign and the level of the bias. Then, any bias can be supported in the equilibrium. This leads to the following proposition.

Proposition 1 *In the one-period model there exists an equilibrium with zero recommendation bias. The profit of firm A in this equilibrium is higher than in any other equilibrium.*

When the bias is zero, all the users choose the content that fits them the best. Hence, their expected utility is maximized and so is the subscription price. In the Nash bargaining solution, firm A gets a weighted sum of its profit when both contents are hosted and its profit when only its own content is hosted, which is its outside option. Firm A's profit in each case is the relevant subscription price, which is equal to users' expected utility, from the full surplus extraction. As the utility when only content A is available does not depend on the bias, zero bias maximizes the profit of firm A.

3 Equilibrium Biases

We now consider the main two-period model. By backward induction, we start by analyzing the second period in Section 3.1. We then move to the analysis of the first period in Section 3.2.

3.1 Period Two

Let b_t^* denote the equilibrium bias for period $t = 1, 2$. Since the second period is the last period of the game, the same reasoning applies as in the benchmark one-period model above: There is an equilibrium with zero second-period bias, $b_2 = 0$, and it maximizes the profit of firm A. Indeed, since the second-period subscription price has been paid, firm A is indifferent with respect to the sign and magnitude of bias b_2. Recommendations with zero bias $b_2 = 0$ provide the best match between content and users' tastes and hence maximize the expected utility of users—which is equal to the subscription price. This argument holds for any first-period bias b_1^* which is the next lemma.

Lemma 1 *Consider an equilibrium with biases b_1^* and b_2^*. Then, there exists another equilibrium $b_1^{*\prime}$ and $b_2^{*\prime}$ with the same first-period bias $b_1^{*\prime} = b_1^*$ and $b_2^{*\prime} = 0$ and it yields a higher profit for firm A.*

As firm B gets a share of the surplus in period 2, it then also gets a higher profit in an equilibrium with biases b_1^* and $b_2^* = 0$ than in any other equilibrium with biases b_1^* and $b_2^* \neq 0$. As firm A always extracts the full surplus from the users, they always get zero surplus. Thus, an equilibrium with biases b_1^* and $b_2^* = 0$ Pareto dominates all other equilibria with b_1^* and $b_2^* \neq 0$. In what follows we will focus on equilibrium with $b_2^* = 0$.

3.2 Period One

We now analyze firm A's incentives to bias its recommendations in the first period. We assume in this section that users follow the recommendation and consider their incentives to do so in Section 3.3. Let b_1^e denote users' expectation about the first-period bias. When firm A chooses some first-period bias b_1, this does not affect its first-period profit since fee f_1 and subscription price p_1 have been set and users have made their subscription decisions. The bias affects, however, its second-period profit. Indeed, the second-period fee f_2 is the solution of the following Nash bargaining problem

$$\max_{f_2} \left(U_{AB,2}(b_1, b_1^e) - f_2 - U_{A,2}(b_1, b_1^e)\right)^{\alpha} f_2^{1-\alpha},$$

where $U_{AB,2}(b_1, b_1^e)$ is the sum of the users' expected utilities when the bargaining succeeds and both contents are available and $U_{A,2}(b_1, b_1^e)$ is the sum of the users' expected utilities when the bargaining fails and only firm A's content is available. Because of full surplus extraction, $U_{AB,2}(b_1, b_1^e)$ is equal to firm A's revenue from subscription prices. The solution generates the following second-period profit of firm A:

$$\Pi_2^A(b_1, b_1^e) = \alpha U_{AB,2}(b_1, b_1^e) + (1-\alpha)\, U_{A,2}(b_1, b_1^e). \tag{2}$$

From (2), the effect of the actual bias b_1 can be decomposed into two effects as follows:

$$\frac{\partial \Pi_2^A(b_1, b_1^e)}{\partial b_1} = \alpha \underbrace{\frac{\partial U_{AB,2}(b_1, b_1^e)}{\partial b_1}}_{\text{perceived surplus effect}} + (1-\alpha) \underbrace{\frac{\partial U_{A,2}(b_1^e, b_1)}{\partial b_1}}_{\text{outside option effect}} \tag{3}$$

The *outside-option* effect—that is, firm A wants to raise its outside option—is standard whenever the surplus is divided according to the Nash bargaining solution.[7] The *perceived-surplus* effect means that firm A wants to raise the surplus consumers expect to obtain when both contents are available. As the perceived surplus depends on the recommendation bias which users do not observe, it can be different from the true surplus. The effect comes into play because the two contents are generically of different quality, $v_A \neq v_B$. Then, raising the perceived surplus requires firm A to recommend more often the ex ante superior content.[8] The weights of the two effects are $1-\alpha$ and α, respectively. A higher bargaining power of the platform α makes it invest less into its outside option and more into the overall surplus perceived by the users.

The next proposition characterizes the equilibrium first-period bias and is the main result of the paper; Figure 1 illustrates Proposition 2.[9]

---

[7] See footnote 3 for the references related to the outside-option effect.

[8] See footnote 3 for the references related to the perceived-surplus effect.

[9] We focus on the equilibrium in which the first-period recommendation is informative at least for some parameter values. There is always a "babbling" equilibrium in which firm A uses an extreme first-period bias and the users ignore the recommendation.

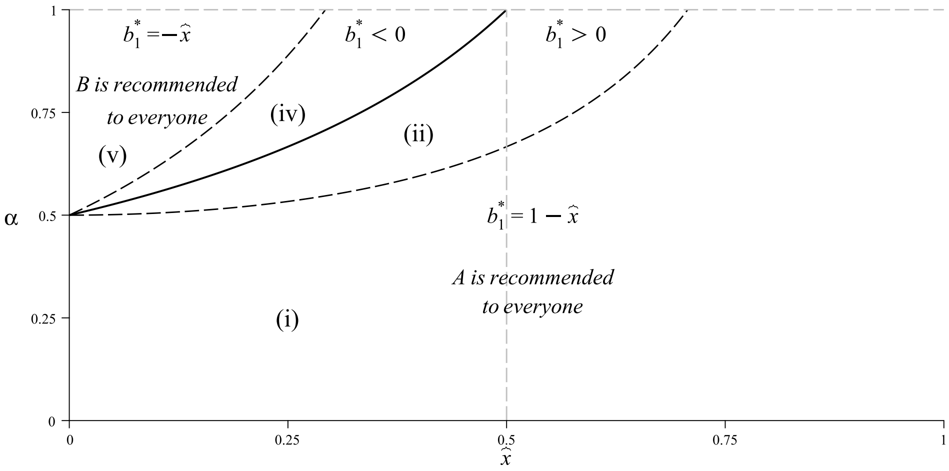

Figure 1: The first-period equilibrium bias b_1^* as a function of $\hat{x}$ and α. The numbers (i), (ii), (iv) and (v) correspond to the cases in Proposition 2. The solid line corresponds to case (iii) (and $b_1^* = 0$) of the same proposition.

Proposition 2 *The equilibrium bias in the first-period b_1^* is characterized as follows:*

- *(i)* $b_1^* = 1 - \hat{x} > 0$ *if* $\alpha \leq \frac{1}{2}\frac{1}{1-\hat{x}^2}$*;*
- *(ii)* $b_1^* = \frac{1-2\alpha(1-\hat{x})}{2\alpha-1}\hat{x} > 0$ *if* $\alpha \in \left[\frac{1}{2}\frac{1}{1-\hat{x}^2}, \frac{1}{2}\frac{1}{1-\hat{x}}\right)$*;*
- *(iii)* $b_1^* = 0$ *if* $\alpha = \frac{1}{2}\frac{1}{1-\hat{x}}$*;*
- *(iv)* $b_1^* = (1-\hat{x})(1-2\alpha(1-\hat{x})) < 0$ *if* $\alpha \in \left(\frac{1}{2}\frac{1}{1-\hat{x}}, \frac{1}{2}\frac{1}{(1-\hat{x})^2}\right]$*;*
- *(v)* $b_1^* = -\hat{x} < 0$ *if* $\alpha \geq \frac{1}{2}\frac{1}{(1-\hat{x})^2}$*.*

b_1^* *is (weakly) decreasing in* α *and increasing in* $\hat{x}$*.*

We start by analyzing the outside-option effect. $U_{A,2}$ in (2) is given as follows:

$$\begin{aligned} U_{A,2}(b_1, b_1^e) &= (\hat{x} + b_1)\left(v_A - \frac{\tau}{2}(\hat{x} + b_1^e)\right) + (1 - (\hat{x} + b_1))\left(v_A - \frac{\tau}{2}(1 + \hat{x} + b_1^e)\right) \\ &= v_A - \frac{\tau}{2}(1 + b_1^e - b_1). \end{aligned} \tag{4}$$

The first term comes from the users that were recommended content A in the first period. There are $\hat{x} + b_1$ of them and they think that their expected distance to A is $(\hat{x} + b_1^e)/2$. The second term, of a similar structure, comes from the users that were recommended content B in the first period.

Therefore, the outside-option effect is given by:

$$\frac{\partial U_{A,2}(b_1^e, b_1)}{\partial b_1} = \frac{\tau}{2}(>0). \tag{5}$$

The sign of the effect is always positive since, for any users' conjecture about the first-period bias b_1^e, biasing recommendation even more towards its own content A makes more users believe that content A is their better match, increasing their expected utility if only content A is available. Therefore, when $\alpha = 0$ (i.e., when the perceived-surplus effect is absent), firm A introduces the maximum bias and recommend A to all the users, which generates an extreme form of self-preferencing. Note that the outside-option effect has a constant size of $\frac{\tau}{2}$ since the difference between two conditional expectations—that is, between the average location of users recommended A and that of those recommended B—is always 1/2 for the uniform distribution. This implies that only A is recommended when α is positive but small (see Figure 1). In this case, even though there also exists the perceived-surplus effect, it is dominated by the outside-option effect.

We now turn to the perceived-surplus effect. Compute $U_{AB,2}$ as

$$U_{AB,2}(b_1, b_1^e) = (\hat{x} + b_1)U_{AB,2}(b_1^e|r_1 = A) + (1 - (\hat{x} + b_1))U_{AB,2}(b_1^e|r_1 = B), \tag{6}$$

where $U_{AB,2}(b_1^e|r_1 = A)$ (respectively, $U_{AB,2}(b_1^e|r_1 = B)$) represents a consumer's expected utility from the second period upon receiving recommendation A (respectively, B) in the first period. With probability $\hat{x} + b_1$, users are recommended content A in the first period and with probability $1 - (\hat{x} + b_1)$, they are recommended content B in the first period.

Crucially, both $U_{AB,2}(b_1^e|r_1 = A)$ and $U_{AB,2}(b_1^e|r_1 = B)$ depend only on the expected bias b_1^e, not on the actual bias b_1. Hence, $U_{AB,2}(b_1, b_1^e)$ in (6) is affine in b_1. The perceived-surplus effect is then given by:

$$\frac{\partial U_{AB,2}(b_1, b_1^e)}{\partial b_1} = U_{AB,2}(b_1^e|r_1 = A) - U_{AB,2}(b_1^e|r_1 = B). \tag{7}$$

The perceived-surplus effect is equal to the difference between users' expected utility from receiving recommendation A and the utility from receiving recommendation B and it is a function of only the expected bias b_1^e.

Consider now the case of $\alpha = 1$. Since there is only the perceived-surplus effect, see (3), in an interior equilibrium the platform has to be indifferent with respect to bias b_1, which implies $U_{AB,2}(b_1^e|r_1 = A) = U_{AB,2}(b_1^e|r_1 = B)$. Given the symmetric distribution of users between the two contents, this is only possible if the users expect that the better content is recommended more often, thus increasing their average distance to that content and, hence, decreasing their expected utility from it. Since in the equilibrium the expected and the actual biases have to be equal, this implies that $b_1^* > 0$ if and only if $v_A > v_B$, that is, if $\hat{x} > 1/2$, see Figure 1. The larger the difference between the two qualities, a higher bias in favor of the better quality content is needed to equalize the expected utilities from the two contents. Thus, the absolute value of b_1^* increases as $\hat{x}$ moves further from 1/2. When the difference between the two qualities is large enough, that is, $\hat{x}$ is close enough to zero or one, equalizing the two utilities is impossible since the support of the users' location distribution is finite. Then, the platform uses an extreme bias by always recommending the better content—which of course makes this recommendation uninformative.

For intermediate α, the two effects work simultaneously. Since the outside-option effect is constant as explained above, all comparative statics is driven by the perceived-surplus effect. As $\hat{x}$ increases, the perceived-surplus effect increases, which in turn raises the equilibrium bias. When α increases—which is the weight of the perceived-surplus effect—the equilibrium bias b_1^* decreases since this effect is smaller than the outside-option effect (when both have the same sign) or has a sign opposite to that of the outside-option effect.

3.3 Incentives to follow the recommendation

Consider now users' incentives to follow the recommendation in period one. Start with the case of a positive first-period equilibrium bias $b_1^* > 0$. Users will follow the recommendation whenever firm A recommends B. Suppose now that firm A recommends A, which occurs for $x \in [0, \hat{x} + b_1^*]$. In this case, the expected utility from following the recommendation is higher than the expected utility from ignoring it and consuming B if

$$v_A - \tau \frac{\hat{x} + b_1^*}{2} \geq v_B - \tau \left(1 - \frac{\hat{x} + b_1^*}{2}\right),$$

which is equivalent to

$$\hat{x} \geq b_1^*. \tag{8}$$

Analyzing similarly the case of a negative first-period bias, we obtain the following result.

Proposition 3 *Given the first-period recommendation bias characterized in Proposition*

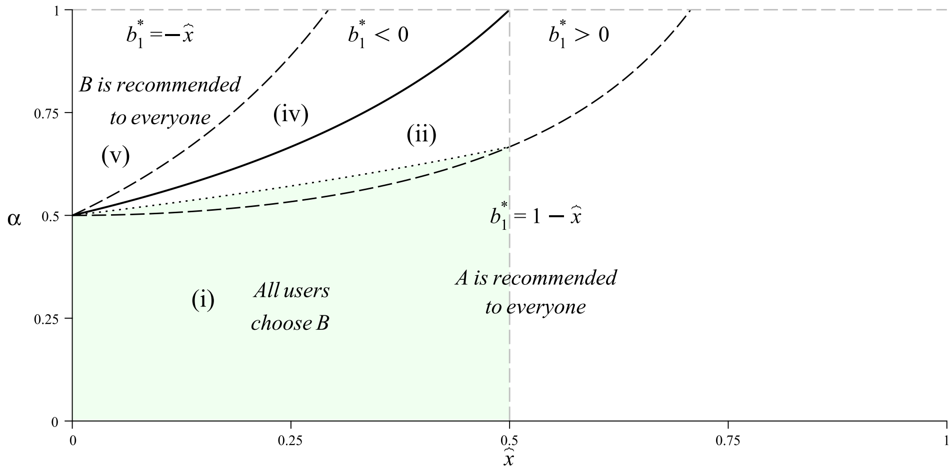

Figure 2: The first-period equilibrium bias b_1^* as a function of $\hat{x}$ and α. In the colored region all users ignore the recommendation and consume B. The numbers (i), (ii), (iv) and (v) correspond to the cases in Proposition 2. The solid line corresponds to case (iii) (and $b_1^* = 0$) of the same proposition.

2, users always follow the recommendation to consume B. They do not follow the recommendation to consume A if $\hat{x} < \frac{1}{2}$ and $\alpha < \frac{1}{2-\hat{x}}$ and follow it otherwise.

Figure 2 illustrates Proposition 3. For intuition, suppose first that the recommendation is completely uninformative—the recommendation is always to consume either content A or B. The users then choose to consume the ex ante better content, that is, content A if $v_A > v_B$ or, equivalently, $\hat{x} > \frac{1}{2}$, and content B otherwise. When the bias is positive but not maximum, recommendation A is informative but inflated. A lower $\hat{x}$, that is, worse quality of A relative to B, and a lower α, that is, a higher weight of the outside-option effect, make the recommendation more inflated and at some point users are better off with the ex ante better quality. Hence, for $\hat{x} < \frac{1}{2}$ and α low enough, users disregard the recommendation A and consume the ex ante better content B. Note that the opposite never happens: Users never disregard recommendation B when it is inflated, i.e., when the bias is negative. Indeed, the negative bias reflects the fact that content B is better since it requires the negative perceived-surplus effect. Then, recommendation B, either completely uninformative if the bias is at the maximum or somewhat informative if it is not, (weakly) increases the expected utility from B which is ex ante better; the users then go with it.

When the users do not follow the recommendation and consume the ex ante better quality, this means that there is no information transmission in the equilibrium. The first-period bias b_1^* characterized in Proposition 2 is still the equilibrium one but it might not be unique. For example, for $\hat{x} < \frac{1}{2}$ and $\alpha \in \left(\frac{1}{2}\frac{1}{1-\hat{x}^2}, \frac{1}{2-\hat{x}}\right)$, that is, in the lower part of region (ii) with $b_1^* > 0$ of Proposition 2 (see Figure 2), any higher bias is also an equilibrium one since the users will still ignore this even less informative recommendation. These multiple equilibria are all payoff-equivalent since no information is transmitted and the users make the same choice.

3.4 The effects of the bias

Let us discuss the effects of the equilibrium bias on profits and consumer gross surplus. As in other signal jamming models, e.g., career concerns of Holmström (1999), in the second period there is no effect—users know that there is no bias there, and with their correct conjecture of the first-period bias they correctly compute their expected utility from the second-period consumption. In the first-period, however, there is a negative effect of the bias—some consumers are misled and hence, their expected utility is lower than without the bias. Then, they pay a lower subscription price, implying that both firms obtain lower profits. Firm A thus would like to refrain from using the bias but it does not have a commitment power for it.

4 Conclusion

This paper studies the incentives of a platform that distributes its own content and third-party content to bias its recommendations to users about which content to consume. The contract between the platform and the content provider is assumed to be a fixed payment decided in the beginning of each period via Nash bargaining. In a static setting, the platform has no incentives to bias the recommendations whereas in a dynamic setting, the platform has such incentives which are shaped by two effects, the outside-option effect and the perceived-surplus effect. While the former one is always in favor of biasing recommendation towards its own content, the direction of the latter one depends on the ex ante relative quality of the contents. Therefore, both self-preferencing and third-party preferencing are possible in the equilibrium, depending on the content qualities and the bargaining weights.

Appendix

Proof of Proposition 1 As written in the text, since the subscription price is paid before the bias is set, firm A is indifferent with respect to the sign and magnitude of the bias. Hence, any bias can be supported in an equilibrium. In particular, there exists an equilibrium with zero bias.

Let us now show that in the equilibrium with zero bias the profit of firm A is maximized. In an equilibrium with bias b, the fee f that firm A pays to firm B is obtained as the solution of the following Nash bargaining problem

$$\max_f \left(U_{AB}(b) - f - U_A\right)^{\alpha} f^{1-\alpha},$$

where $U_{AB}(b)$ is the users' expected utility when the bargaining succeeds and $U_A = v_A - \tau/2$ is the users' expected utility when the bargaining fails and only firm A's content is available. Because of full surplus extraction, $U_{AB}(b)$ (respectively, U_A) is equal to firm A's revenue from subscription prices. The resulting profit of firm A is then

$$\Pi^A(b) = \alpha U_{AB}(b) + (1-\alpha)\left(v_A - \tau/2\right). \tag{9}$$

The users' expected utility $U_{AB}(b)$ is[10]

$$U_{AB}(b) = (\hat{x}+b)\left(v_A - \tau\frac{\hat{x}+b}{2}\right) + (1-(\hat{x}+b))\left(v_B - \tau\Big(1 - \frac{1+\hat{x}+b}{2}\Big)\right),$$

where the first term comes from the users to whom content A is recommended and the second term comes from the users to whom content B is recommended. In the equilibrium the users have the correct beliefs $b^e = b$. Then,

$$\frac{\partial}{\partial b}\Pi^A(b) = \alpha\frac{\partial}{\partial b}U_{AB}(b) = -2\alpha\tau b.$$

Hence, $\Pi^A(b)$ is hump-shaped with the maximum at $b=0$. ■

Proof of Lemma 1 Similarly to (9), in an equilibrium with biases b_1^* and b_2 profit of firm A in the second period is

$$\Pi_2^A(b_1^*, b_2) = \alpha U_{AB,2}(b_1^*, b_2) + (1-\alpha)\, U_{A,2}(b_1^*).$$

---

[10] This expression assumes that users follow the recommendation. If they do not, their utility is $\max\{v_A, v_B\} - \frac{1}{2}\tau$. For a sufficiently small bias they will follow the recommendation.

If the bargaining fails, only content A is available and there are no recommendations. Hence, $U_{A,2}(b_1^*)$ does not depend on b_2.

Let us compute $U_{AB,2}(b_1^*, b_2)$. Suppose first that $b_1^* \geq b_2$. Then, there are three groups of users: Those in $[0, \hat{x} + b_2]$ who were recommended A in both periods; users in $[\hat{x} + b_2, \hat{x} + b_1^*]$ were recommended A in period 1 and B in period 2, and those in $[\hat{x} + b_1^*, 1]$ that were recommended B in both periods. Hence,[11]

$$\begin{aligned} U_{AB,2}(b_1^*, b_2) =& (\hat{x} + b_2)\Big(v_A - \tau \frac{\hat{x} + b_2}{2}\Big) + (b_1^* - b_2)\Big(v_B - \tau\Big(1 - \hat{x} - \frac{b_1^* + b_2}{2}\Big)\Big) \\ &+ (1 - (\hat{x} + b_1^*))\Big(v_B - \tau \frac{1 - (\hat{x} + b_1^*)}{2}\Big). \end{aligned} \tag{10}$$

If $b_1^* \leq b_2$, the three groups of users are those in $[0, \hat{x} + b_1^*]$ were recommended A in both periods; users in $[\hat{x} + b_1^*, \hat{x} + b_2]$ were recommended B in period 1 and A in period 2, and those in $[\hat{x} + b_2, 1]$ were recommended B in both periods. Hence,

$$\begin{aligned} U_{AB,2}(b_1^*, b_2) =& (\hat{x} + b_1^*)\Big(v_A - \tau \frac{\hat{x} + b_1^*}{2}\Big) + (b_2 - b_1^*)\Big(v_A - \tau\Big(\hat{x} + \frac{b_2 + b_1^*}{2}\Big)\Big) \\ &+ (1 - (\hat{x} + b_2))\Big(v_B - \tau \frac{1 - (\hat{x} + b_2)}{2}\Big). \end{aligned} \tag{11}$$

Then, using (1) for $\hat{x}$, in both cases (10) and (11) $\frac{\partial}{\partial b_2} U_{AB,2}(b_1^*, b_2) = -2\tau b_2$. Hence, $U_{AB,2}(b_1^*, b_2)$ is hump-shaped with the maximum at $b_2 = 0$ and the same then applies to $\Pi_2^A(b_1^*, b_2)$. ■

Proof of Proposition 2 Expression (4) gives $U_{A,2}$. Compute now $U_{AB,2}$. We need to distinguish two cases, $b_1^e \gtrless 0$. Start with $b_1^e \geq 0$:

$$\begin{aligned} U_{AB,2}(b_1, b_1^e | b_1^e \geq 0) =& (\hat{x} + b_1)\left[\frac{\hat{x}}{\hat{x} + b_1^e}\Big(v_A - \tau \frac{\hat{x}}{2}\Big) + \frac{b_1^e}{\hat{x} + b_1^e}\Big(v_B - \tau\Big(1 - \Big(\hat{x} + \frac{1}{2} b_1^e\Big)\Big)\Big)\right] \\ &+ (1 - (\hat{x} + b_1))\Big(v_B - \tau\Big(1 - \frac{1 + \hat{x} + b_1^e}{2}\Big)\Big). \end{aligned}$$

With probability $\hat{x} + b_1$, users are recommended content A in the first period. These users think that they are located in $[0, \hat{x} + b_1^e]$ and that in the second period they will be recommended A if their location is in $[0, \hat{x}]$ and B if it is in $[\hat{x}, \hat{x} + b_1^e]$. With probability $1 - (\hat{x} + b_1)$, users are recommended content B in the first period. These users think that

---

[11] Expressions (10) and (11) assume that users follow the recommendation which is the case if b_2 is small enough.

they are located in $[\hat{x} + b_1^e, 1]$ and that they will be recommended B again in the second period.

Since both $U_{A,2}(b_1, b_1^e)$ and $U_{AB,2}(b_1, b_1^e | b_1^e \geq 0)$ are affine functions of b_1, $\Pi_2^A(b_1, b_1^e)$ in (2) is affine too. Then, in an equilibrium in which the bias is not extreme, firm A has to be indifferent with respect to the bias b_1 (see also footnote 9). Solving $\frac{\partial \Pi_2^A(b_1, b_1^e)}{\partial b_1} = 0$ yields $b_1^e = \frac{1-2\alpha(1-\hat{x})}{2\alpha - 1}\hat{x}$ and it is (weakly) positive if $\alpha \leq \frac{1}{2}\frac{1}{1-\hat{x}} > \frac{1}{2}$. For $\alpha < \frac{1}{2}\frac{1}{1-\hat{x}^2}$, $\frac{1-2\alpha(1-\hat{x})}{2\alpha-1}\hat{x} > 1 - \hat{x}$ and, hence $b_1^* = 1 - \hat{x}$ in this case. If $\alpha \leq \frac{1}{2}$, then $\frac{\partial \Pi_2^A(b_1, b_1^e)}{\partial b_1} > 0$ for any $b_1^e \geq 0$ and hence, in the equilibrium all the users are recommended content A in the first period, that is, $b_1^* = 1 - \hat{x}$.

Suppose now that $b_1^e \leq 0$. Then,

$$U_{AB,2}(b_1, b_1^e | b_1^e \leq 0) = (\hat{x} + b_1)\Big(v_A - \tau\frac{\hat{x} + b_1^e}{2}\Big)$$
$$+(1 - (\hat{x} + b_1))\left[\frac{-b_1^e}{1 - (\hat{x} + b_1^e)}\Big(v_A - \tau\Big(\hat{x} + \frac{b_1^e}{2}\Big)\Big) + \frac{1 - \hat{x}}{1 - (\hat{x} + b_1^e)}\Big(v_B - \tau\Big(1 - \frac{1 + \hat{x}}{2}\Big)\Big)\right].$$

With probability $\hat{x} + b_1$ users are recommended content A in the first period. They think that in the second period they will be also recommended A. With probability $1 - (\hat{x} + b_1)$ users are recommended content B in the first period; then, they think that they will be recommended content A in the second period if their location is in $[\hat{x} + b_1^e, \hat{x}]$ and content B if their location is in $[\hat{x}, 1]$.

As in the case $b_1^e \geq 0$ above, $\Pi_2^A(b_1, b_1^e)$ is affine in b_1. Hence, firm A has to be indifferent with respect to bias b_1 in the equilibrium. Solving $\frac{\partial \Pi_2^A(b_1, b_1^e)}{\partial b_1} = 0$ yields $b_1^e = (1 - \hat{x})(1 - 2\alpha(1 - \hat{x}))$ and it is (weakly) negative if $\alpha \geq \frac{1}{2}\frac{1}{1-\hat{x}} > \frac{1}{2}$. Finally, $(1 - \hat{x})(1 - 2\alpha(1 - \hat{x})) < -\hat{x}$ if $\alpha \geq \frac{1}{2}\frac{1}{(1-\hat{x})^2}$ in which case $b_1^e = -\hat{x}$, that is, all the users are recommended content B in the first period.

Combining the two cases, we get the statement of the proposition. ■

Proof of Proposition 3

Consider first the case of $b_1^* > 0$. The users follow recommendation to consume B. The condition for following recommendation to consume A is given by (8): $\hat{x} \geq b_1^*$. If $b_1^* = 1 - \hat{x}$ (case (i) of Proposition 2), then (8) yields $\hat{x} \geq \frac{1}{2}$. If $b_1^* = \frac{1-2\alpha(1-\hat{x})}{2\alpha-1}\hat{x}$ (case (ii) of Proposition 2), then (8) yields $\alpha \geq \frac{1}{2-\hat{x}}$. Combining these conditions with relevant ranges for α from the two cases and taking its complement to $[0, 1]^2$ results in the region specified in the proposition.

Consider now the case of $b_1^* < 0$. Then, the users follow the recommendation to consume A. Suppose now that the recommendation is to consume B, which happens for $x \in [\hat{x} + b_1^*, 1]$. In this case, the expected utility from following the recommendation is higher than the utility from consuming A if

$$v_A - \tau \frac{1 + \hat{x} + b_1^*}{2} \leq v_B - \tau \left(1 - \frac{1 + \hat{x} + b_1^*}{2}\right)$$

which is equivalent to

$$|b_1^*| \leq 1 - \hat{x}. \tag{12}$$

From cases (iv) and (v) of Proposition 2, the necessary condition for $b_1^* < 0$ is $\hat{x} < \frac{1}{2}$ while the maximum bias is in case (v), $b_1^* = -\hat{x}$. Hence, (12) is always satisfied and the users always follow the recommendation if $b_1^* < 0$. ■

References

Abdollahpouri, Himan, Masoud Mansoury, Robin Burke, Bamshad Mobasher, and Edward Malthouse. 2021. “User-centered Evaluation of Popularity Bias in Recommender Systems.” *UMAP ’21*, 119–129.

Aguiar, Luis, Joel Waldfogel, and Sarah Waldfogel. 2021. “Playlisting favorites: Measuring platform bias in the music industry.” *International Journal of Industrial Organization*, 78: 102765.

Bourreau, Marc, and Germain Gaudin. 2022. “Streaming platform and strategic recommendation bias.” *Journal of Economics & Management Strategy*, 31(1): 25–47.

Brandenburger, Adam, and Ben Polak. 1996. “When Managers Cover Their Posteriors: Making the Decisions the Market Wants to See.” *RAND Journal of Economics*, 27(3): 523–541.

Brin, Sergey, and Larry Page. 1998. “The Anatomy of a Large-Scale Hypertextual Web Search Engine.” Seventh International World-Wide Web Conference (WWW 1998), April 14-18, 1998, Brisbane, Australia.

Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello. 2026. “Artificial Intelligence, Algorithmic Recommendations and Competition.” `https://ssrn.com/abstract=4448010`.

Celdir, Musa Eren, Soo-Haeng Cho, and Elina H. Hwang. 2024. “Popularity Bias in Online Dating Platforms: Theory and Empirical Evidence.” *Manufacturing & Service Operations Management*, 26(2): 537–553.

Che, Yeon-Koo, Wouter Dessein, and Navin Kartik. 2013. “Pandering to Persuade.” *American Economic Review*, 103(1): 47–79.

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. 2019. “A model of biased intermediation.” *RAND Journal of Economics*, 50(4): 854–882.

European Commission. 2017. “Antitrust: Commission fines Google €2.42 billion for abusing dominance as search engine by giving illegal advantage to own comparison shopping service - Factsheet." 27 June 2017, available at `http://europa.eu/rapid/press-release_MEMO-17-1785_en.htm`.

European Commission. 2022. "Antitrust: Commission accepts commitments by Amazon barring it from using marketplace seller data, and ensuring equal access to Buy Box and Prime." 20 December 2022, available at `https://ec.europa.eu/commission/presscorner/api/files/document/print/en/ip_22_7777/IP_22_7777_EN.pdf`.

Farronato, Chiara, Andrey Fradkin, and Alexander MacKay. 2023. "Self-Preferencing at Amazon: Evidence from Search Rankings." *AEA Papers and Proceedings*, 113: 239–43.

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.

Gomez-Uribe, Carlos A., and Neil Hunt. 2015. "The Netflix Recommender System: Algorithms, Business Value, and Innovation." *ACM Transactions on Management Information Systems*, 6(4): Article 13.

Grossman, Sanford J., and Oliver D. Hart. 1986. "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration." *Journal of Political Economy*, 94(4): 691–719.

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(Supplement C): 48–60.

Hagiu, Andrei, Tat-How Teh, and Julian Wright. 2022. "Should platforms be allowed to sell on their own marketplaces?" *RAND Journal of Economics*, 53(2): 297–327.

Hart, Oliver, and John Moore. 1990. "Property Rights and the Nature of the Firm." *Journal of Political Economy*, 98(6): 1119–1158.

Heidhues, Paul, and Johan Lagerlöf. 2003. "Hiding information in electoral competition." *Games and Economic Behavior*, 42(1): 48–74.

Holmström, Bengt. 1999. "Managerial Incentive Problems: A Dynamic Perspective." *Review of Economic Studies*, 66(1): 169–182.

Kittaka, Yuta, Susumu Sato, and Yusuke Zennyo. 2023. "Self-preferencing by platforms: A literature review." *Japan and the World Economy*, 66: 101191.

Klein, Benjamin, Robert G. Crawford, and Armen A. Alchian. 1978. "Vertical Integration, Appropriable Rents, and the Competitive Contracting Process." *Journal of Law & Economics*, 21(2): 297–326.

Lee, Dokyun, and Kartik Hosanagar. 2019. "How Do Recommender Systems Affect Sales Diversity? A Cross-Category Investigation via Randomized Field Experiment." *Information Systems Research*, 30(1): 239–259.

Motta, Massimo. 2023. "Self-preferencing and foreclosure in digital markets: Theories of harm for abuse cases." *International Journal of Industrial Organization*, 90: 102974.

Nicas, Jack, and Keith Collins. 2019. "How Apple's App Store Is under Threat from Antitrust Regulators." *The New York Times*. https://www.nytimes.com/interactive/2019/09/09/technology/apple-app-store-competition.html, accessed: 2024-05-25.

Rey, Patrick, and Jean Tirole. 2007. "A Primer on Foreclosure." In *Handbook of Industrial Organization*. Vol. 3, , ed. M. Armstrong and R. Porter, 2145–2220. Elsevier.

Vestager, Margrethe. 2017. Speech at the Bundeskartellamt 18th Conference on Competition, Berlin, 16 March 2017, available at https://ec.europa.eu/commission/commissioners/2014-2019/vestager/announcements/bundeskartellamt-18th-conference-competition-berlin-16-march-2017_en.