Product Recommendations with Match Externalities

Preston Boorsma*

September 24, 2023

Abstract

In designing product recommendation systems, platforms face a tradeoff between providing "safe" recommendations (repetitive or familiar recommendations) that they know consumers are likely to entertain and "discovery" recommendations (recommendations generated from collaborative filtering or other discovery-oriented algorithms) that consumers may be less likely to entertain yet potentially provide more value than safe recommendations. Committing to provide discovery recommendations when safe recommendations are of little value increases buyers' willingness to pay for a recommendation service but may also lower the total number of successful consumer-producer matches created through the service. I show how match externalities to third-party advertisers can tip the scale of this tradeoff towards providing more safe recommendations to prioritize volume of successful matches over consumer value from participation. Although some consumers would be better-off with new product discovery, the platform provides them with safe recommendations to increase its revenue from third-party advertisers. The platform compensates consumers for their loss in service value through a lower participation fee. While consumer recommendation efficiency deteriorates with the existence of third-party advertisers, consumer surplus increases. When faced with a decision to invest in recommendation system "allure" (the match-likelihood of discovery recommendations) or recommendation system "suitability" (the match-conditional expected value of discovery recommendations), the platform prefers to invest in allure in most cases.

* Department of Economics, Michigan State University, Email: boorsmap@msu.edu

1 Introduction

Online platforms like streaming service providers (e.g., YouTube; Netflix; Spotify) and marketplaces (e.g., Amazon.com; Booking.com) provide consumers access to effectively limitless product options that, due to search frictions, would provide them little incremental value without any platform recommendation system to orient their search. Platform recommendation systems are limited by historical consumer-producer interactions and other available data. This introduces a tradeoff for platforms in their recommendation decisions. On one hand, platforms can use historical individual-consumer-producer interactions data to provide repetitive or similar product recommendations to consumers. If individual consumer preferences do not change too much over time, then these are "safe" recommendations for the platform in that a consumer is likely to be willing to interact with these recommended producers. On the other hand, platforms can provide experimental recommendations (e.g., through collaborative filtering) to promote new product discovery by consumers. While these "discovery" recommendations may provide more value to some consumers compared to their safe recommendations, the platform faces more uncertainty about consumers' willingness to interact with them. This paper formally studies this safe-experimental recommendation tradeoff faced by platforms, along with the implications it has for consumers.

A recent study by Chen et al. (2023) clearly demonstrates this tradeoff on a music streaming platform. The authors experimentally implemented a relatively more discovery-oriented recommendation system, and they show that it increased some users' consumption diversity but decreased overall consumption on the platform. Different platforms implement systems on differing sides of the recommendation tradeoff. For example, Ricks and McCrosky (2022) shows that although YouTube provides tools for consumers to provide feedback on videos, user feedback often does not impact future video recommendations. From complementary survey data, the authors document that consumers on YouTube must resort to other tactics like strategically adjusting viewing activities to influence their video recommendations. This suggests that YouTube's recommendation system may prioritize volume of engagement over value from engagement. In contrast, Netflix places high emphasis on providing recommendations that maximize consumer value—it has even released data and hosted a public prize contest in attempt to improve its ability to predict individual customer content preferences (not match-likelihoods; Thompson, 2008). A key difference between these two platforms is their revenue structures. YouTube is mainly advertisement-funded, while Netflix is subscription-funded.[1] Due to third-party advertisers on YouTube, a

[1] Ricks and McCrosky (2022) do not gather the share of participants in their study who have premium (ad-free) and free (ad-present) YouTube accounts. When much of their data was collected in 2021, less than 3.5%

video view yields a positive externality from the consumer-producer interaction of advertiser access to a consumer's attention. The value of this view-externality depends not on how much the consumer likes the video but only whether the consumer watches the video. This could explain why YouTube prioritizes engagement over consumer value relative to Netflix. I formalize this intuition and other results by modeling platforms' recommendation efficiency tradeoff under varying levels of match-externalities.

To study this issue, I build a model in which a platform provides consumers with access to many producers for a participation fee. Due to search frictions, consumers rely on the platform to provide them with a product recommendation. For each consumer, the platform has a "safe" producer, which the consumer is highly likely to "match" with if recommended, and a "discovery" producer, which the consumer is less likely to "match" with if recommended. Consumers have knowledge about the value they would receive from safe recommendations due to their similarity to unmodelled previous interactions, and they may communicate their preference for safe recommendations to the platform. Consumers are heterogenous in their valuation for a safe recommendation. The consumers and platform only know distributional features of the match-conditional value they would receive from a discovery recommendation. Thus, consumers with low value safe recommendations prefer to receive a discovery recommendation even though it is less likely that such an experimental producer will be a successful match. Third-party advertisers value access to successful consumer-producer matches, and the platform sells this access for a fee. The key friction here is that third-party advertiser preferences only depend on volume of engagement by consumers—whether consumers experience successful matches—while consumers with low safe recommendation values would prefer a low match-likelihood discovery recommendation over their safe recommendation. I allow the value of the advertising market to vary. The model encompasses many platform service structures. For example, media streaming and social media platforms provide advertisement space for which demand increases in volume of engagement by consumers but does not depend on how much consumers value their interactions. Similarly, online marketplaces may sell advertisements in complementary product spaces once transactions are completed. After formally developing the model, I thoroughly relate it and its assumptions to the online video streaming platform YouTube to demonstrate its applicability to a specific example platform.

When analyzing the model, I first characterize the platform's choice of recommendation strategy and consumer- and advertiser-participation fees. As the value of the advertising market increases, the platform provides discovery recommendations to consumers who prefer

of all YouTube users were estimated to be premium subscribers, suggesting that their results are largely driven by the ad-present version of YouTube (https://www.statista.com/statistics/1261865/youtube-premium-subscribers/; https://www.statista.com/forecasts/1144088/youtube-users-in-the-world).

them less often in order to increase overall match volume and increase the advertiser participation fee. Consumers are charged lower participation fees in order to sustain high consumer participation levels. In order to sustain high participation levels from consumers with low safe recommendation valuations, the platform cannot charge consumers for the heterogenous values they receive from safe recommendations. Thus, even though consumers experience lower quality recommendations, consumer surplus generally increases with advertiser value. This reveals a tradeoff between the prevalent "premium" and "free" platform subscription structures for consumers. While they may receive better recommendations and greater aggregate value with advertisement-free services compared to advertisement-funded services, the increase in participation fee incurred may significantly offset any value generation from improved recommendation efficiency.

Next, I study the platform's incentives to invest in discovery recommendation quality. Recommendation quality is made up of two components: match-likelihood or "allure" of discovery recommendations and match-conditional expected value or "suitability" of discovery recommendations. Increasing recommendation allure or suitability increases consumer value from discovery recommendations in similar ways, but increasing allure also produces a positive advertiser demand effect by increasing the match-likelihood for all discovery recommendations. I demonstrate that when advertiser value is sufficiently high, the platform always prefers to invest in allure compared to suitability when investing in recommendation quality. Further, at any fixed recommendation quality level, the platform always prefers to tradeoff recommendation suitability in favor of recommendation allure. These results provide an explanation for repetitive or consumption-narrowing product recommendation "filter bubbles" often observed in practice.

The rest of the paper proceeds as follows. Section 2 reviews the relevant literature on recommendation systems. Section 3 builds the model, illustrates it with a motivating example, and discusses its assumptions in depth. Section 4 analyzes the model and develops the main results, with all proofs provided in an Appendix. Section 5 concludes.

2 Literature Review

Much of the literature on product recommendation systems studies the implementation and effects of specific types of recommendation systems. Anderson et al. (2020) provide evidence that Spotify's recommendation algorithms reduce individual users' listening diversity.[2] The authors show that nonincreasing listening diversity over time is associated

[2] Anderson et al. (2020)'s findings would suggest that Spotify's subscription funded service prioritizes volume of engagement over value from engagement through its personalized recommendations, which does not align

with higher levels of recommendation-based listening and lower levels of user-driven listening, while increasing listening diversity over time is associated with lower levels of recommendation-based listening and higher levels of user-driven listening. Holtz et al. (2020) show that personalized podcast recommendations on Spotify based on listening history increase engagement but decrease listening diversity compared to popularity-based recommendations. Anderson et al. (2020) and Holtz et al. (2020) provide evidence for a common view that personalized recommendation systems can create "filter bubbles" through which consumers repeatedly see similar products in recommendation cycles that reduce individual consumption diversity and new product discovery (Pariser, 2011). Other authors present evidence that recommendation systems can instead increase consumption diversity. Brynjolfsson et al. (2011) and Oestreicher-Singer and Sundararajan (2012) study online retail industries and demonstrate that recommendation systems can redistribute demand from "popular" products to "niche" products, suggesting that recommendations can increase overall consumption diversity. Fleder and Hosanagar (2009) reconcile these contradicting implementation effects and shows that recommendation systems may either increase or decrease individual and aggregate consumption diversity, depending on the system being implemented. Chen et al. (2023) and Holtz et al. (2020) further illustrate this point by showing how a recommendation algorithm can effectively be adjusted to further promote music listening diversity. Lee and Wright (2023) also explore the effects of recommendation systems. They show how collaborative filtering can generate value for consumers depending on training data, the number of recommendations considered, and the complexity of prediction inputs. Many authors focus on specific consumer preference learning processes, which are relevant to implementation of certain recommendation systems (e.g., Kremer et al., 2014; Che and Hörner, 2018; Feng et al., 2022).

Rather than looking into the implementation and effects of specific recommendation systems, the focus of this paper is to look into a consumption diversity tradeoff faced in the design of recommendation systems. Fewer papers have taken this approach. Fleder and Hosanagar (2009) suggest that recommendation algorithms that promote discovery may be better for firms and consumers, depending on the environment. However, they make this claim only by considering two "popular" recommendation algorithms and comparing welfare effects between the two systems. Chen et al. (2023) empirically study the effects of more

with the predictions of this paper. However, the underlying tradeoff persists, and this tension may occur due to suboptimal recommendation practices by Spotify. Chen et al. (2023) demonstrate through a field experiment that a more discovery-oriented recommendation system on a similar music streaming platform lowered overall consumption but increased consumption diversity for some users. Even with decreased volume of engagement, the platform ultimately adopted the more discovery-oriented recommendation system. This suggests that although Spotify prioritizes volume of engagement, it may be profitable to shift its recommendations to further promote discovery and increase value from engagement.

diverse music recommendations on listeners' behaviors on a music platform. They show that a more discovery-oriented recommendation system increased active users' consumption diversity without reducing their consumption levels, but the discovery-oriented recommendation system had no effect on overall consumption diversity and decreased overall consumption when considering all users. Importantly, the accuracy with which users' preferences could be predicted determined the ability of recommendations to increase consumption diversity. Chen et al. (2023) make it clear that promotion of new product discovery through recommendations can come at a cost of reducing consumption for consumers who receive poor recommendations. Studying actual implementation of recommendation systems is a complex topic in itself and is not the objective of this paper. However, papers that do so clearly demonstrate that platforms have a significant degree of choice in the outcomes of recommendation systems, especially along the lines of consumption diversity (Chen et al., 2023; Fleder and Hosanagar, 2009; Anderson et al., 2020; Holtz et al., 2020). To study recommendation system design rather than implementation, I take as given that a platform can implement certain outcomes from its recommendation system and abstract away from any specifics on how these outcomes are implemented.

With a similar approach, numerous papers study incentives for platforms to bias recommendations in cases of sponsored search. These authors mainly focus on cases in which platforms have incentives to divert search in order to increase (decrease) transfers paid by (to) producers (de Corniére and Taylor, 2014; Bourreau and Gaudin, 2021; Hagiu and Jullien, 2011) or to preference a platform's own products over its competitors' products (Hagiu et al., 2022; Aridor and Gonçalves, 2022; Zou and Zhou, 2023). In contrast, I abstract away from sponsored or prominent search incentivized by producers and focus on a form of recommendation bias that arises when third-party advertisers benefit from successful consumer-producer matches.

This paper takes several prevalent platform features as given to focus on the role of recommendations systems, while other work focuses on these features. This paper takes as given wide consumer-producer access provision through platforms, and is most relevant to such platforms, while Hagiu and Wright (2023b) study the strategic discoverability choice concerning how much consumer access a platform should provide to participating producers. I also take as given a platform's ability to monetize consumer-producer access provision, whereas other papers study producer disintermediation tactics and platform policy responses (Edelman and Wright, 2015; Boorsma, 2023; Hagiu and Wright, 2023a).

3 A Model of Product Recommendations

3.1 Model

There are four types of players: producers, consumers, third-party advertisers, and a monopoly platform that facilitates discovery between consumers and producers and provides advertisement space through successful consumer-producer matches.

Producers. A unit mass of producers indexed by $j \in J = [0,1]$ produce a horizontally differentiated product. Each producer has zero marginal cost and zero outside option. Each producer is thus willing to accept a zero lump-sum transfer from the platform to supply the market. These simplifying assumptions allow us to ignore the supply side of the market to focus on recommendations.

Consumers. A unit mass of consumers indexed by $i \in I = [0,1]$ have unit demand for the product. Each consumer i has valuation $m_{ij}u_{ij} \in [0,1]$ for producer j's product, where $m_{ij} \in \{0,1\}$ is a binary "match" indicator that determines whether consumer i is willing to try j's product, and u_{ij} is consumer i's match-conditional valuation for j. Consumers have a homogenous outside option value of u_0, and the valuations $m_{ij}u_{ij}$ follow some known distribution with expectation less than u_0. Because of this, consumers will not sample a product randomly and depend on the platform for a recommendation. Following Lee and Wright (2023), consumers may only sample one product. This could be due to lack of time or attention or due to the good being an experience good. I omit the consumer index i in most notation.

Advertisers. A unit mass of advertisers indexed by $k \in K = [0,1]$ benefit from visibility in consumer-producer matches. Let n_m denote the mass of successfully matched interactions between producers and consumers. Each advertiser k has valuation $\beta_k n_m^\alpha$ for advertisement through the platform, where $\alpha \in (0,1)$ is a constant and the β_k follow a uniform distribution G_β over $\left[0, \bar{\beta}\right]$.[3] I omit the advertiser index k in most notation.

Platform. A monopoly platform facilitates discovery between consumers and producers through product recommendations to consumers. The platform can provide one of two types of recommendations to consumers: "safe" and "discovery." From the collection of consumer data (e.g., previous consumption data), the platform has a "safe" recommended producer $S_i \in J$ for each consumer i, where $m_{S_i} = 1$ with certainty.[4] The safe recommendation

[3] The results do not qualitatively rely on the chosen functional form for advertisers' valuations, but diminishing marginal returns to the mass of successful matches introduces a smooth tradeoff between safe and discovery recommendations where one recommendation type is not always better than the other for all consumers.

[4] The safe recommendation need not have unit match-likelihood. All results hold as long as a safe recommendation has a higher match-likelihood than a discovery recommendation.

valuations u_{S_i} follow a uniform distribution G_S over $[0,1]$. The platform may also provide a "discovery" recommended producer $D_i \in J$, where $m_{D_i} = 1$ with fixed probability ρ. The discovery recommendation valuations u_{D_i} follow a known distribution over $[0,1]$ with expectation μ. The value ρ may be interpreted as the "allure" of discovery recommendations, and the expected match quality μ may be interpreted as the match-conditional "suitability" of the discovery recommendations. Finally, the platform has two revenue sources. It may charge a participation fee f_c to consumers, and it may charge a participation fee f_a to advertisers.

Consumers and the platform have better information about the values of each u_{S_i} because the S_i are safe recommendations in the sense that the same or similar recommendations and/or interactions have previously occurred and been observed. I assume that each consumer i knows the value of u_{S_i}, but the platform only knows the distribution of the u_{S_i} and that each $m_{S_i} = 1$. With a consumer-invariant participation fee, the most the platform can learn from self-revelation of the u_{S_i} by consumers is equivalent to allowing each consumer to send a signal $s \in \{S, D\}$ to the platform specifying their desired recommendation type—safe or discovery.[5] This matches commonly observed tools on platforms that allow consumers to provide feedback on product recommendations. A recommendation strategy for the platform is a function $\sigma\colon \{S, D\} \to [0,1]$, where $\sigma(s)$ specifies the probability with which the platform recommends D_i when it receives signal s.

The timing of the game is as follows. First, the platform sets its fees and publicly announces its recommendation strategy. Next, consumers realize their values of the u_{S_i}, decide whether to participate, and send signals to the platform. Advertisers decide whether to participate. Finally, consumers receive recommendations according to the platform's strategy, matches and discovery valuations are realized, and all payoffs are realized. Without consequence to any results, the realization of the u_{S_i} and signal sending could occur first in the game. Realization of the u_{S_i} occurs through some unmodelled learning through past interactions—a recommendation system implementation feature studied by other authors, that I abstract from.

[5] A consumer's expected value is maximized with recommendation D if and only if $u_S < \rho\mu$. Consider a platform strategy that recommends D with probability $\sigma(u_S)$ upon consumer revelation of u_S. Without a negative consequence due to the invariance of f_a, any participating consumer with $u_S < \rho\mu_D$ would reveal $\arg\max \sigma_F(u_S)$, and any participating consumer with $u_S \geq \rho\mu$ would reveal $\arg\min \sigma_F(u_S)$. So truthful revelation requires that $\sigma(\cdot)$ has a range made up of no more than two values. I consider agent-invariant fees in the main specification to match frequently practiced fee structures.

3.2 Illustrative Example and Discussion of Assumptions

The model encompasses many different platform services, but it is useful to illustrate the main features through a motivating example platform, YouTube. I provide further discussion of the modeling assumptions along the way.

Players. YouTube is an online video sharing service, and all three players in addition to the platform are present in the market. Video creators (producers) upload content to the platform and consumers watch videos through the platform. Third-party advertisers benefit from access to consumer attention when consumers watch creators' videos. This benefit depends on whether consumers watch creators' videos, not how much consumers value watching creators' videos.

A need for recommendations. Consumers face inordinate video choices—over 500 hours of videos were freely uploaded to YouTube every minute as of June 2022.[6] Supposing there was no YouTube recommendation system in place to orient consumer search, additional video uploads would yield virtually no marginal benefit to consumers due to search frictions. Briefly departing from the model, YouTube would have an incentive to increase upload costs for creators to improve expected quality of content from search on the platform and increase consumer value or engagement. Such high volume of daily video uploads should not be observed without a recommendation system. Instead, by procuring video recommendations for consumers, YouTube makes video consumption feasible, and additional video uploads can improve consumer value or engagement even with so many videos already available on the platform.

Safe and discovery recommendations. Consumers generally watch many videos through YouTube over time, so the platform accumulates historical consumer viewing behavior on each consumer. As long as consumer preferences do not change too much, past viewing behavior is predictive of future viewing of repetitive or similar content—safe recommendations have higher match-likelihoods than experimental recommendations ($\rho < 1$). However, because of the vast and continually increasing availability of content on YouTube, as well as differing amounts of data the platform has on individual consumers, product discovery through collaborative filtering or other experimental recommendation algorithms may yield higher expected value for some consumers depending on their viewing history and consequent safe recommendations ($\mu > 0$).

Single-product sampling. In the model, consumers only consider one product recommendation (as in Lee and Wright, 2023). This is a simplifying assumption, but it may still be justified. First, if evaluating the recommended/prominent video is costless (e.g., as in Armstrong et al., 2009; Zou and Zhou, 2023) or less costly than evaluating another video,

[6] https://www.statista.com/statistics/259477/hours-of-video-uploaded-to-youtube-every-minute/

and if subsequent evaluation costs are high relative to value from consumption, then a consumer may only be willing to consider one recommendation. Alternatively, consumers may experience a lack of attention beyond the first recommendation. For example, if YouTube videos are considered a "distraction" for consumers, then their outside option may increase with the number of products they evaluate. Finally, online videos may be considered an experience good. Consumers may not have time to sample another video once they have started evaluating a recommended video.[7]

In considering single-product sampling, I have assumed that consumers know their safe recommendation valuations but cannot watch the safe recommendation if they receive their discovery recommendations, even if the discovery recommendations do not produce successful matches. The lack of attention or experience good justifications for single-product sampling can also rationalize this assumption. It may also be the case that consumers do not do not know which videos are similar to those they have watched in the past, but they recall the value they received from their past experiences. Without consequence to any results, we can also think of the realization of the u_{S_i} and message-sending interactions taking place in the first stage of the game. This view more closely resembles what happens on YouTube, where consumers can provide feedback (e.g., video "likes" or "dislikes") while they are watching a video for future recommendation input. In that case, consumers may not be able to find a repeat or similar video when it is not recommended to them if they have poor recall of the previous learning and feedback stage. Regardless of any mapping from the simplified single-product sampling setting to reality, the key element introduced with single-product sampling that drives results in the model is that consumers generate less advertisement revenue for the platform when they receive more recommendations that are not successful matches.

Recommendation quality. Quality of discovery recommendations are characterized in the model by their match-likelihood or "allure" ρ and their match-conditional expected value or "suitability" μ. On YouTube, a video description must be alluring to entice consumers to click on them to watch. If a platform knows more about what catches the interest of consumers, then it can supply more alluring recommendations. An alluring video may be either suitable or unsuitable. While many niches of video may catch the attention of any given consumer, some video niches may provide the consumer with greater satisfaction from actually watching the videos compared to others. Overall discovery recommendation quality $\rho\mu$ depends on the platform's data and investment in its algorithms.

[7] In the case of experience goods, we must assume that consumers can evaluate products for their match success $m_j \in \{0, 1\}$ before giving any attention to third-party advertisers. If a producer is not a successful match, then a consumer stops watching the video before seeing any advertisements.

Participation fees. On YouTube, video creators may freely upload content to the platform. Some producers are paid through advertisement-revenue sharing with the platform, but I abstract away from this feature. Consumers may freely watch videos on the platform ($f_c = 0$), potentially due to an effective non-negative consumer price constraint (considered below). Third-party advertisers pay for visibility when consumers watch videos ($f_a > 0$). As shown below, the platform's optimal recommendation strategy depends on the participation fees it charges. It is interesting to note that there is evidence that YouTube recommendations are inefficient from a consumer's point of view, in line with the outcomes developed in the analysis below. With no consumer and positive advertiser participation fees, the platform does not always have an incentive to provide consumers with their desired recommendation types. Ricks and McCrosky (2022) shows that YouTube often does not provide recommendations in line with consumer-generated input.

No "Freemium" Menu. YouTube and other streaming and media services offer a "freemium" consumer service menu in practice, in which consumers may choose between a high-cost (premium) advertisement-free and a free advertisement-present service option. The model here captures recommendation tradeoff incentives within either service option. Relevant platforms likely adopt a menu of service options to price discriminate along heterogenous preferences for certain service features, like consumers' nuisance costs from exposure to advertisements (Sato, 2019; Jeon et al., 2022). I study within-service recommendation design incentives and do not allow for multiple service options because platforms likely do not price discriminate along consumers' preferences for familiar content (u_{S_i}). If these preferences change over time and a platform price discriminated based on them, then we should observe consumers often switching back and forth between premium and free options. This seems unrealistic, and heterogenous nuisance costs and overall service quality preferences likely motivate any dual service menu provision in practice.

4 Analysis

4.1 Equilibrium Characterization

I solve for the optimal platform strategy $\sigma(s)$ and participation fees f_c and f_a. I focus on an equilibrium satisfying the natural condition $\sigma(D) \geq \sigma(S)$. In such a case a participating consumer will signal $s = D$ iff $u_S < \rho\mu$. Compared to any positive value of $\sigma(S)$, the platform strictly prefers $\sigma(S) = 0$. On the consumer-side, setting $\sigma(S) = 0$ does not change D-signalling consumer payoffs but increases S-signalling payoffs, allowing for a weakly higher f_c holding all else equal. On the advertiser-side, setting $\sigma(S) = 0$ increases the total number of successful matches created because discovery recommendations have lower match-likelihoods than safe recommendations ($\rho < 1$); this allows for a weakly higher f_a holding all else equal. The platform thus sets $\sigma(S) = 0$, and, from an S-signalling consumer's perspective, platform recommendations are efficient.

Lemma 1. The platform always recommends the safe recommendation S to consumers who prefer it: $\sigma(S) = 0$.

Taking $\sigma(S) = 0$ as given, denote $\sigma \equiv \sigma(D)$ to simplify notation. Employing Lemma 1, an S-signalling consumer (with $u_S \geq \rho\mu$) will participate iff $f_c \leq u_S - u_0$. A D-signalling consumer (with $u_S < \rho\mu$) will participate iff

$$f_c \leq (1-\sigma)u_S + \sigma\rho\mu - u_0 \Leftrightarrow u_S \geq \frac{f_c + u_0 - \sigma\rho\mu}{1-\sigma}.$$

Then the mass of consumers who participate at the fee level f_c and discovery probability σ is given by

$$n_c = \begin{cases} 1 - G_S\left(\dfrac{f_c + u_0 - \sigma\rho\mu}{1-\sigma}\right), & f_c \leq \rho\mu - u_0, \\ 1 - G_S(f_c + u_0), & f_c > \rho\mu - u_0. \end{cases}$$

The mass of consumers with successful matches at the fee level f_c and discovery probability σ is given by

$$n_m = \begin{cases} [1-\sigma(1-\rho)]\left[G_S(\rho\mu) - G_S\left(\dfrac{f_c + u_0 - \sigma\rho\mu}{1-\sigma}\right)\right] + 1 - G_S(\rho\mu), & f_c \leq \rho\mu - u_0, \\ 1 - G_S(f_c + u_0), & f_c > \rho\mu - u_0. \end{cases}$$

Note that $f_c \leq \rho\mu - u_0$ is the condition necessary for any D-signalling consumer to participate. If $f_c > \rho\mu - u_0$, then participation is not profitable for D-signalling consumers even if they always get their preferred recommendation ($\sigma = 1$).

An advertiser will participate iff $f_a \leq \beta n_m^\alpha$. The mass of advertisers who participate at the fee level f_a and discovery probability σ is given by

$$n_a = 1 - G_\beta\left(\frac{f_a}{n_m^\alpha}\right).$$

Note that there is only an indirect network effect from consumer participation to advertiser value, and consumers are unaffected by advertiser participation. Having derived consumer and advertiser demand, we can write the platform's profit as

$$\Pi(f_c, f_a, \sigma) = f_c n_c + f_a n_a.$$

For $f_c > \rho\mu - u_0$, D-signalling consumers' expected benefit from discovery recommendations do not cover their costs to participate, only S-signalling consumers may be willing to participate, and platform profit does not depend on σ. In such a case, the platform will set the standard monopoly platform market participation fee f_c (e.g., Armstrong, 2006) to be the opportunity cost of an additional increase in the consumer participation fee (a loss in advertiser revenue through the indirect network effect) adjusted upwards by a factor related to the elasticity of consumer participation, and it will employ standard monopoly pricing for f_a due to a lack of any indirect network effect from advertisers to consumers. The focus of this paper is on the tradeoff between safe and discovery recommendations, but in this case discovery recommendations are of too poor quality to warrant any induced participation by those who would benefit from them. In what remains, I assume that the discovery recommendation quality $\rho\mu$ is sufficiently large so that the platform optimally sets $f_c \leq \rho\mu - u_0$ and at least some consumers who prefer discovery recommendations participate. A sufficient condition is $\rho\mu \geq \frac{1}{4}(1 + 2u_0 + u_0^2)$, which ensures that $f_c \leq \rho\mu - u_0$ is optimal even with no market for advertising ($\bar{\beta} = 0$).

Taking $f_c \leq \rho\mu - u_0$ as given, consider any platform choice (f_c, σ). All consumers participate iff $f_c \leq \sigma\rho\mu - u_0$. If $f_c < \sigma\rho\mu - u_0$, then the platform could increase f_c to improve consumer-side profits without any effect on advertiser-side participation or profits. If $f_c > \sigma\rho\mu - u_0$, then an increase in σ has two opposing effects on the number of successful matches n_m. First, an increase in σ improves D-signalling consumers' payoffs from participating because they more often get their desired recommendation; thus, an increase in σ draws more D-signalling consumers into the market, a $\big(1 - \sigma(1-\rho)\big)$-fraction of whom end up with successful matches. Second, an increase in σ decreases the number of successful matches from D-signalling consumers who participated before the change in σ because these consumers get the lower match-likelihood discovery recommendation more often than before. Lemma 2 verifies that with uniformly distributed safe valuations u_S, the positive first effect always dominates the negative second effect.

Lemma 2. The consumer participation fee satisfies $f_c = \sigma\rho\mu - u_0$, and all consumers participate.

Lemma 2 can also be interpreted as follows. The platform induces full consumer participation by setting $f_c = \sigma\rho\mu - u_0$. Since consumer participation is already full at this fee level, lowering σ in attempt to further increase n_m is a necessary action to further improve advertiser-side profits through consumer-side mechanisms. Keeping f_c fixed, consumer participation lost from a reduction in σ has a greater negative effect on n_m than matches gained from remaining consumer participants who more often receive the high match-likelihood safe recommendation. Thus, if the platform looks to increase n_m to improve advertiser-side profits, it must do so by decreasing σ *and* decreasing f_c to maintain full consumer participation.

From Lemmas 1 and 2 and the demand equations derived above, the platform's profit simplifies to

$$\begin{aligned}\Pi(f_a, \sigma) &= \sigma\rho\mu - u_0 + f_a \left[1 - G_\beta\left(\frac{f_a}{n_m^\alpha}\right)\right] \\ n_m &= 1 - \sigma(1-\rho)G_S(\rho\mu).\end{aligned}$$

It is then straightforward to derive the optimal advertiser participation fee and discovery recommendation probability, summarized below.

Proposition 1. The platform always recommends S to consumers who prefer it, and it recommends D to consumers who prefer it with the discovery recommendation probability

$$\sigma = \begin{cases} 1, & \bar{\beta} < \dfrac{4}{\alpha(1-\rho)}[1-(1-\rho)\rho\mu]^{1-\alpha}, \\ \dfrac{1}{(1-\rho)\rho\mu}\left[1 - \left(\dfrac{\bar{\beta}\alpha(1-\rho)}{4}\right)^{\frac{1}{1-\alpha}}\right], & \bar{\beta} \in \dfrac{4}{\alpha(1-\rho)}[[1-(1-\rho)\rho\mu]^{1-\alpha}, 1], \\ 0, & \bar{\beta} > \dfrac{4}{\alpha(1-\rho)}. \end{cases}$$

The optimal participation fees are given by $f_a = \frac{1}{2}\bar{\beta}[1-\sigma(1-\rho)\rho\mu]^\alpha$ and $f_c = \sigma\rho\mu - u_0$. Furthermore:

(i) In response to an increase in advertiser value $\bar{\beta}$, the discovery recommendation probability σ decreases, the advertiser-participation fee f_a increases, and the consumer-participation fee f_c decreases.

(ii) In response to an increase in discovery recommendation allure ρ, the discovery recommendation probability σ increases, the advertiser-participation fee f_a either increases or decreases, and the consumer-participation fee f_c increases.

(iii) In response to an increase in discovery recommendation suitability μ, the discovery recommendation probability σ decreases, the advertiser-participation fee f_a decreases, and the consumer-participation fee f_c increases.

The proposition fully characterizes the platform's recommendation decisions and fee levels. Below I further explore the roles of advertiser value $\bar{\beta}$ and the components of discovery recommendation quality $\rho\mu$ in equilibrium outcomes, using discussion of Proposition 1 as a starting point.

4.2 Advertiser Value and Equilibrium Outcomes

Figure 1 illustrates the platform's equilibrium choices as a function of advertiser value $\bar{\beta}$ for specific values of the other exogenous parameters (u_0, ρ, μ, α).

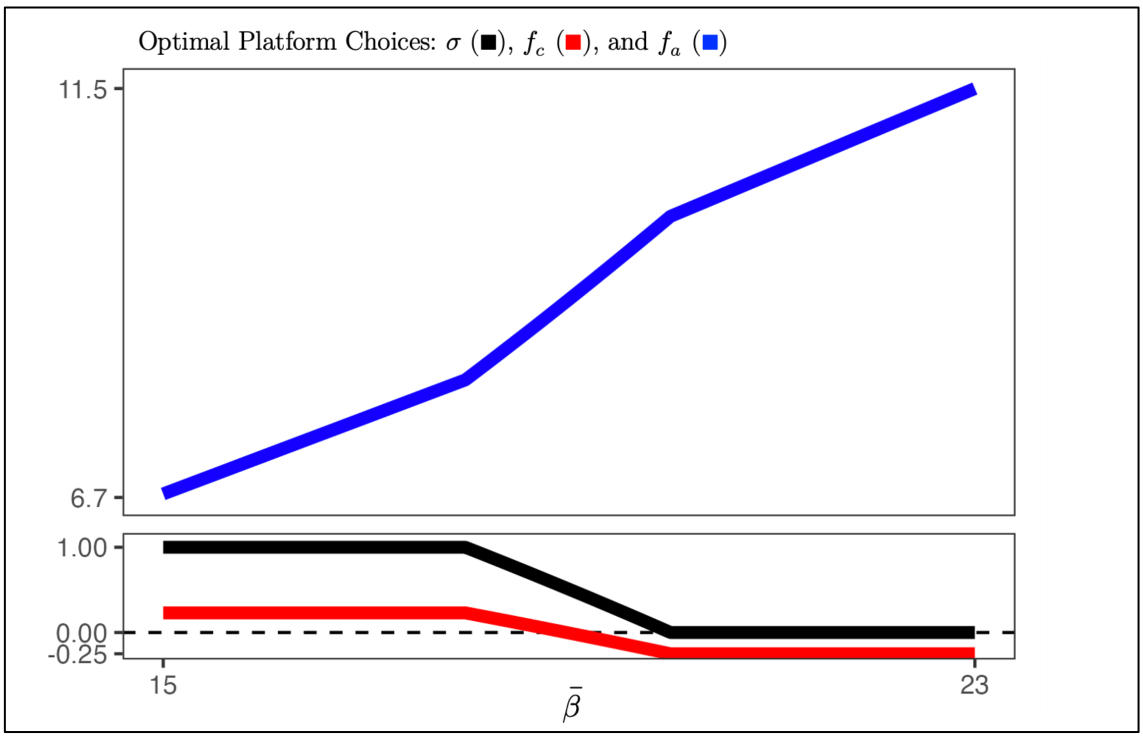

Figure 1
Parameters: $u_0 = 0.25, \rho = 0.6, \mu = 0.8, \alpha = 0.5$

Intuitively, advertiser participation fees increase with advertisers' value from match access $\bar{\beta}$. Consumer participation fees decrease with $\bar{\beta}$, and they become negative and approach $-u_0$ when advertisers value match-access sufficiently highly. Decreasing consumer participation fees f_a are driven by deteriorated consumer recommendation efficiency. As advertisers more highly value access to consumer-producer matches through an increase in

$\bar{\beta}$, recommendation efficiency for D-signalling consumers decreases through a decrease in σ; that is, as consumer-producer match externalities increase, consumer value generation decreases to preserve high match-likelihoods. More consumers do not receive their preferred recommendations because the platform's main revenue source shifts from consumers to advertisers. Interestingly, however, while consumer *value* decreases with higher match externalities, consumer *surplus* increases.

Proposition 2. Consumer recommendation efficiency is decreasing in advertiser value $\bar{\beta}$, while consumer surplus is increasing in advertiser value $\bar{\beta}$. Platform profit is increasing in advertiser value $\bar{\beta}$.

Consumers are better-off with higher match externalities despite lowered consumer value generation because the platform compensates consumers with lower fees in order to sustain full participation. S-signalling consumers receive the same recommendations for all $\bar{\beta}$, but they face a lower participation fee for higher levels of $\bar{\beta}$. They are better-off with higher match externalities. D-signalling consumers receive their desired recommendations less often with higher $\bar{\beta}$, but they also face a lower participation fee for higher levels of $\bar{\beta}$. A D-signalling consumer earns a surplus of $u_0 + (1-\sigma)u_S$, which decreases in σ and consequently increases in $\bar{\beta}$. D-signalling consumers are fully charged for the discovery recommendations they expect to receive but do not pay for the safe recommendations they expect to receive. They are better-off with higher match externalities. To further explore this result, it is useful to compare the extreme case $\bar{\beta} = 0$ with higher levels of $\bar{\beta}$. When $\bar{\beta} = 0$, the platform sets $\sigma = 1$ and D-signalling consumers all receive the homogenous ex ante value from participation $\rho\mu$. The platform can fully extract this value net of the outside option while sustaining full consumer participation. When $\bar{\beta}$ is sufficiently large, the platform sets $\sigma < 1$ and D-signalling consumers receive heterogenous values from participation $\sigma\rho\mu + (1-\sigma)u_S$, which depend on u_S. The platform cannot extract any of the heterogenous participation value $(1-\sigma)u_S$ without losing participation from consumers with the lowest values of u_S. The platform thus forfeits some surplus to D-signalling consumers to sustain full participation. Proposition 2 has significant practical relevance. It reveals a tradeoff between the prevalent "premium" and "free" platform subscription structures for consumers. While they may receive better recommendations and greater aggregate value with advertisement-free services compared to advertisement-funded services, the increase in participation fee incurred may significantly offset value generation from improved recommendation efficiency.

4.2 Discovery Recommendation Quality and Equilibrium Outcomes

Proposition 2 derives from analysis of the effects of advertiser value $\bar{\beta}$ on equilibrium outcomes through Proposition 1(i). I now further consider the effects of discovery recommendation allure ρ and suitability μ on equilibrium outcomes through Proposition 1(ii) and Proposition 1(iii). It is first interesting to note that even though an increase in allure ρ or suitability μ increases overall discovery recommendation quality $\rho\mu$ in similar ways, the equilibrium discovery probability σ increases in ρ but decreases in μ. Whether these effects are strict and consequent effects on platform fees charged depend on advertiser value for match access. I compare the effects of increases in ρ and μ for low, intermediate, and high levels of $\bar{\beta}$ in turn, which correspond to cases when platform discovery recommendations fully favor consumer demands ($\sigma = 0$), balance consumer and advertiser demands ($0 < \sigma < 1$), and fully favor advertiser demands ($\sigma = 1$), respectively.

First consider low levels of advertiser value ($\bar{\beta} < \frac{4}{\alpha(1-\rho)}[1-(1-\rho)\rho\mu]^{1-\alpha}$). The platform always gives consumers their preferred recommendations ($\sigma = 1$) and prioritizes consumer value generation over volume of engagement as a revenue source. An increase in ρ or μ has a larger positive effect on consumer compared to advertiser demand, so the platform continues to prioritize consumer value generation and does not change its recommendation strategy from $\sigma = 1$ after such a change. Because consumer value increases, the platform raises the consumer participation fee f_c. It lowers the advertiser fee f_a in response to an increase in μ, as n_m decreases because more consumers prefer D and σ remains unchanged. An increase in ρ, however, has an indeterminate effect on n_m because it increases the number of D-signalling consumers but also increases the match-likelihood for D-signalling consumers; the platform may either increase or decrease f_a in response to an increase in ρ.

Now consider intermediate levels of advertiser value ($\bar{\beta} \in \frac{4}{\alpha(1-\rho)}[[1-(1-\rho)\rho\mu]^{1-\alpha}, 1]$). The platform sets $\sigma \in (0,1)$ and balances both consumer value generation and volume of engagement as revenue sources. An increase in ρ or μ increases the value consumers receive from the discovery recommendation and also increases the number of consumers that prefer the discovery recommendation. In response to an increase in μ the platform adjusts the discovery recommendation probability σ downwards to offset any negative effect on advertiser demand from an influx in D-signalling consumers and leaves n_m unchanged. Both f_c and f_a remain unchanged. Compared to an increase in μ, an increase in ρ additionally increases the match-likelihood for consumers who receive the discovery recommendation. An increase in ρ thus exhibits an advertiser demand effect that an increase in μ does not exhibit. This positive effect on n_m from increased discovery recommendation match-likelihood counteracts some of the negative effect from increased D-signalling consumers, and the platform increases σ in response to an increase in ρ. The platform increases f_c because consumers receive better and more efficient recommendations. The platform lowers advertiser fees f_a because its revenue source shifts towards consumer value generation enough that n_m decreases.

Finally consider high levels of advertiser value ($\bar{\beta} > \frac{4}{\alpha(1-\rho)}$). The platform never provides consumers with the discovery recommendation ($\sigma = 0$) and prioritizes advertiser value from volume of engagement over consumer value from engagement as a revenue source. An increase in ρ or μ has a larger positive effect on consumer compared to advertiser demand, so the platform may either maintain its low-quality recommendations with $\sigma = 0$ or begin to provide discovery recommendations with $\sigma > 0$. If consumer value indeed increases from the platform providing more discovery recommendations, then the platform charges consumers a higher fee f_c. With any new discovery recommendations provided, the number of successful matches n_m necessarily decreases, and the platform charges advertisers a lower fee f_a.

Throughout the above analysis, we see that all else equal, an increase in allure ρ or suitability μ increases discovery recommendation quality hence consumer value from participation, but an increase in allure ρ does this in addition to having a positive effect on advertiser demand. This is because advertiser demand and the value of the consumer-producer match externality depend only on whether a successful match is made and not on how much value is generated within a successful match. This suggests that the platform may have more incentive to invest in discovery recommendation allure compared to suitability, as the following proposition confirms.

Proposition 3. Platform profit increases in both discovery recommendation allure ρ and suitability μ, but the platform has more incentive to invest in allure compared to suitability with match-externalities in the following two senses:
(i) Fixing discovery recommendation quality $\rho\mu$, the platform always weakly prefers to tradeoff suitability μ in favor of allure ρ. The preference is strict when $\sigma > 0$ and $\bar{\beta} > 0$.
(ii) For all (ρ, μ) there exists $\tilde{\beta} > 0$ such that $\frac{\partial \Pi^*}{\partial \rho} \geq \frac{\partial \Pi^*}{\partial \mu}$ for all $\bar{\beta} \geq \tilde{\beta}$. The inequality holds strictly in a right-open interval containing $\tilde{\beta}$.

Figure 2 illustrates Proposition 3 for specific values of the exogenous parameters $\left(u_0, \alpha, \bar{\beta}\right)$. Optimal platform profits are represented by color on iso-recommendation-quality curves made up of points $(\rho, \mu) \in \{(\rho, \mu) \colon \rho\mu \geq \frac{1}{4}(1 + 2u_0 + u_0^2)\}$ such that recommendation quality $\rho\mu$ is the same for all points on each curve. Observe that platform profit increases in both discovery recommendation allure and suitability—profit increases (color gets darker) when travelling north or east from any point in the (ρ, μ)-plane. In general the platform tends to prefer higher levels of allure ρ compared to suitability μ, as profit is higher (colors are darker) to the right of the 45-degree line, where $\rho > \mu$. To see

Proposition 3(i) in Figure 2, observe that on any iso-recommendation-quality curve, the platform always prefers to tradeoff suitability for allure—profit increases (color gets darker) when travelling counterclockwise on any iso-recommendation-quality curve. Since $\bar{\beta}$ is sufficiently high in Figure 2, Proposition 3(ii) is relevant in that platform profit increases more from an increase in discovery recommendation allure compared to an equivalent increase in suitability—profit increases more (color gets darker) when travelling east compared to north from any point in the (ρ, μ)-plane.

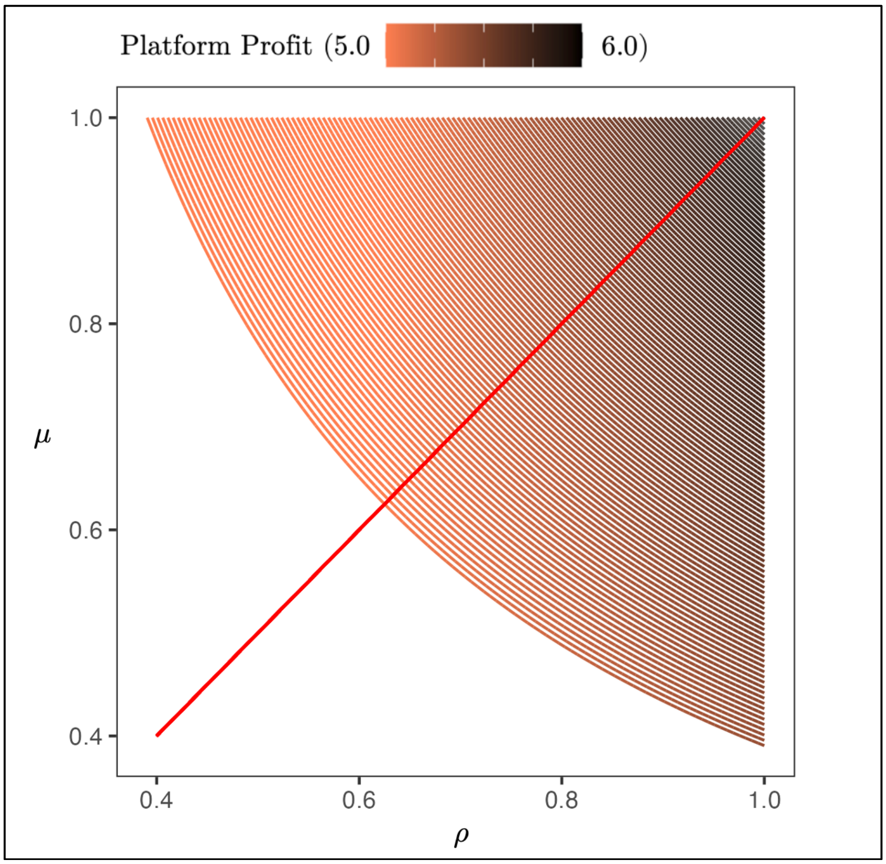

Figure 2
Parameters: $u_0 = 0, \bar{\beta} = 20, \alpha = 0.5$

Note that the existence of values $\bar{\beta}$ for which $\sigma \in (0,1)$ relies on diminishing advertiser returns to consumer-producer matches n_m ($\alpha < 1$). If instead $\alpha = 1$, then the platform always chooses $\sigma \in \{0,1\}$. Even in some cases with $\alpha < 1$, $\sigma \in (0,1)$ on a relatively small interval of $\bar{\beta}$-values. This may seem to make Proposition 3(i) less relevant when $\bar{\beta}$ is sufficiently high so that $\sigma = 0$ and the platform is actually indifferent between investment in ρ and μ.[8] Note, though, that when $\sigma = 0$ the platform optimally sets $f_c = -u_0$. In practice, a platform may face a non-negative consumer price constraint due to adverse selection or opportunistic behaviors by consumers (Amelio and Jullien, 2012; Choi and Jeon, 2021). With a zero unconstrained optimal discovery probability σ, the platform would set $f_c = 0$ to satisfy a non-negative price constraint and choose $\sigma > 0$ to sustain full consumer participation. Thus Proposition 3(i) remains relevant whenever the platform faces a non-negative consumer price constraint and $\sigma = 0$ due to sufficiently high advertiser value for consumer-producer match access. In fact, with a nonzero outside option u_0 and a non-negative consumer price constraint, Proposition 3(i) can be strengthened to state that the platform always strictly prefers to tradeoff suitability μ in favor of allure ρ at any fixed discovery recommendation quality $\rho\mu$ when $\bar{\beta} > 0$.

5 Conclusion

This paper identifies a key tradeoff in the development and implementation of recommendation systems. The quality of product recommendations is limited by historical data. On one hand, historical interactions may be used to provide familiar recommendations to consumers that are "safe" in the sense that consumers are likely to interact with them because they have interacted with them or similar producers in the past. On the other hand, a vast supply of producers makes it unlikely that "safe" recommendations provide the best match for all consumers. Platforms might thus provide "discovery" recommendations that may provide more value to consumers at the risk of a lower successful match-likelihood. I show how consumer-producer match externalities to third party advertisers may incentivize a platform to deteriorate recommendation quality—providing consumers with their undesired recommendation types—to increase profits from advertisers. Despite lower recommendation quality, consumers end up better-off with higher match-externalities because they face lower fees from the platform. To maintain high consumer participation despite lower value from participation, the platform does not charge for poor quality recommendations and forfeits some surplus to consumers. Due to consumer-producer match externalities that depend on whether matches occur and not the value produced in matches, platforms often have higher incentives to invest in more "alluring" recommendation systems that exhibit higher match-likelihoods compared to more "suitable" recommendation systems that provide higher match-conditional expected values to consumers.

[8] When $\sigma = 0$, the platform would not consider any costly investment in recommendation quality $\rho\mu_D$ unless it made enough improvement to support $\sigma > 0$.

References

Amelio, A., & Jullien, B. (2012). Tying and freebies in two-sided markets. *International Journal of Industrial Organization*, *30*(5), 436-446.

Anderson, A., Maystre, L., Anderson, I., Mehrotra, R., & Lalmas, M. (2020, April). Algorithmic effects on the diversity of consumption on spotify. In *Proceedings of The Web Conference 2020* (pp. 2155-2165).

Armstrong, M. (2006). Competition in two-sided markets. *The RAND journal of economics*, *37*(3), 668-691.

Armstrong, M., Vickers, J., & Zhou, J. (2009). Prominence and consumer search. *The RAND Journal of Economics*, *40*(2), 209-233.

Aridor, G., & Gonçalves, D. (2022). Recommenders' originals: The welfare effects of the dual role of platforms as producers and recommender systems. *International Journal of Industrial Organization*, *83*, 102845.

Boorsma, P. (2023). Platform Coherence Policies with a Multiproduct Seller. *Available at SSRN 4531596*.

Bourreau, M., & Gaudin, G. (2022). Streaming platform and strategic recommendation bias. *Journal of Economics & Management Strategy*, *31*(1), 25-47.

Brynjolfsson, E., Hu, Y., & Simester, D. (2011). Goodbye pareto principle, hello long tail: The effect of search costs on the concentration of product sales. *Management science*, *57*(8), 1373-1386.

Che, Y. K., & Hörner, J. (2018). Recommender systems as mechanisms for social learning. *The Quarterly Journal of Economics*, *133*(2), 871-925.

Chen, G., Chan, T., Zhang, D., Liu, S., & Wu, Y. (2023). The Effects of Diversity in Algorithmic Recommendations on Digital Content Consumption: A Field Experiment. *Available at SSRN 4365121*.

Choi, J. P., & Jeon, D. S. (2021). A leverage theory of tying in two-sided markets with nonnegative price constraints. *American Economic Journal: Microeconomics*, *13*(1), 283-337.

de Corniére, A., & Taylor, G. (2014). Integration and search engine bias. *The RAND Journal of Economics*, *45*(3), 576-597.

Edelman, B., & Wright, J. (2015). Price coherence and excessive intermediation. *The Quarterly Journal of Economics*, *130*(3), 1283-1328.

Feng, Y., Caldentey, R., & Ryan, C. T. (2022). Robust learning of consumer preferences. *Operations Research*, *70*(2), 918-962.

Fleder, D., & Hosanagar, K. (2009). Blockbuster culture's next rise or fall: The impact of recommender systems on sales diversity. *Management science*, *55*(5), 697-712.

Ricks, B., & McCrosky, J. (2022). Does This Button Work? Investigating YouTube's Ineffective User Controls. *Mozilla Foundation*.

Hagiu, A., & Jullien, B. (2011). Why do intermediaries divert search?. *The RAND Journal of Economics*, *42*(2), 337-362.

Hagiu, A., Teh, T. H., & Wright, J. (2022). Should platforms be allowed to sell on their own marketplaces?. *The RAND Journal of Economics*, *53*(2), 297-327.

Hagiu, A., & Wright, J. (2023a). Marketplace leakage. *Management Science*.

Hagiu, A., & Wright, J. (2023b). Optimal discoverability on platforms. Working Paper.

Holtz, D., Carterette, B., Chandar, P., Nazari, Z., Cramer, H., & Aral, S. (2020, July). The engagement-diversity connection: Evidence from a field experiment on spotify. In *Proceedings of the 21st ACM Conference on Economics and Computation*(pp. 75-76).

Jeon, D. S., Kim, B. C., & Menicucci, D. (2022). Second-degree price discrimination by a two-sided monopoly platform. *American Economic Journal: Microeconomics*, *14*(2), 322-369.

Kremer, I., Mansour, Y., & Perry, M. (2014). Implementing the "wisdom of the crowd". *Journal of Political Economy*, *122*(5), 988-1012.

Lee, G. and Wright, J. (2023). Recommender systems and the value of user data. Working Paper.

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

Pariser, E. (2011). *The filter bubble: How the new personalized web is changing what we read and how we think*. Penguin.

Sato, S. (2019). Freemium as optimal menu pricing. *International Journal of Industrial Organization*, *63*, 480-510.

Thompson, C. (2008, November 21). If you liked this, you're sure to love that. *The New York Times Magazine*.

Zou, T., & Zhou, B. (2023). Self-preferencing and Search Neutrality in Online Retail Platforms.

Appendix: Proofs

Proof of Lemma 2. If $f_c < \sigma\rho\mu - u_0$, then the consumer-participation fee can be increased without affecting n_c or n_a to increase profit. Suppose $\sigma\rho\mu - u_0 < f_c (\leq \rho\mu - u_0$ by assumption), and notice

$$\begin{aligned}\frac{\partial n_a}{\partial \sigma} &= \alpha g_\beta\left(\frac{f_a}{n_m^\alpha}\right)\frac{f_a}{n_m^{\alpha+1}}\left\{[1-\sigma(1-\rho)]g_S\left(\frac{f_c+u_0-\sigma\rho\mu}{1-\sigma}\right)\frac{\rho\mu-u_0-f_c}{(1-\sigma)^2}\right. \\ &\qquad \left. -(1-\rho)\left[G_S(\rho\mu)-G_S\left(\frac{f_c+u_0-\sigma\rho\mu}{1-\sigma}\right)\right]\right\} \\ &= \alpha\rho g_\beta\left(\frac{f_a}{n_m^\alpha}\right)\frac{f_a}{n_m^{\alpha+1}}\frac{\rho\mu-u_0-f_c}{(1-\sigma)^2} \geq 0.\end{aligned}$$

That is, if $f_c > \sigma\rho\mu - u_0$, then the platform can increase σ with a non-negative effect on advertiser demand. After increasing σ, the platform could increase f_c so as to keep consumer demand unchanged and increase overall profit. Then unless $f_c = \sigma\rho\mu - u_0$, the platform has a profitable deviation, and this must hold in equilibrium. The result implies $n_c = 1$. ■

Proof of Proposition 1: The first order condition for f_a is given by

$$f_a = n_m^\alpha \frac{1-G_\beta\left(\frac{f_a}{n_m^\alpha}\right)}{g_\beta\left(\frac{f_a}{n_m^\alpha}\right)}.$$

Since the platform will clearly set $f_a \in \left[0, \bar{\beta} n_m^\alpha\right]$ this condition yields the unique solution $f_a = \frac{1}{2}\bar{\beta} n_m^\alpha$. The first order condition for σ is given by

$$\rho\mu = \frac{f_a^2}{n_m^{\alpha+1}} \alpha(1-\rho) G_S(\rho\mu) g_\beta \left(\frac{f_a}{n_m^\alpha}\right) \Leftrightarrow \sigma = \frac{1}{(1-\rho)\rho\mu} \left\{1 - \left[\frac{\bar{\beta}\alpha(1-\rho)}{4}\right]^{\frac{1}{1-\alpha}}\right\}.$$

Accounting for probabilistic constraints on σ yields the stated solution. To simplify notation, denote $\bar{\beta}_1 \equiv \frac{4}{\alpha(1-\rho)}[1-(1-\rho)\rho\mu]^{1-\alpha}$ and $\bar{\beta}_2 \equiv \frac{4}{\alpha(1-\rho)}$, where $\left(\bar{\beta}_1, \bar{\beta}_2\right)$ is the interval in which $\sigma \in (0,1)$. Statement (i) follows from direct inspection of the optimal platform choices. It is clear that $\frac{\partial \sigma}{\partial \bar{\beta}} < 0$ and $\frac{\partial f_c}{\partial \bar{\beta}} < 0$ for $\bar{\beta} \in \left[\bar{\beta}_1, \bar{\beta}_2\right]$ and $\frac{\partial \sigma}{\partial \bar{\beta}} = \frac{\partial f_c}{\partial \bar{\beta}} = 0$ otherwise; $\frac{\partial f_a}{\partial \bar{\beta}} > 0$ for all $\bar{\beta}$.

Consider statement (ii). It follows that $\frac{\partial \sigma}{\partial \rho} \geq 0$ from the following computations:

$$\begin{aligned}
\frac{\partial \bar{\beta}_1}{\partial \rho} &= \frac{\bar{\beta}_1}{1-\rho}\left[1 + (1-\alpha)(1-\rho)\mu \frac{2\rho - 1}{1-(1-\rho)\rho\mu}\right] > 0. \\
\frac{\partial \bar{\beta}_2}{\partial \rho} &= \frac{4}{\alpha(1-\rho)^2} > 0, \\
\left.\frac{\partial \sigma}{\partial \rho}\right|_{\bar{\beta} \in (\bar{\beta}_1, \bar{\beta}_2)} &= \frac{1}{(1-\rho)^2\rho\mu}\left\{\frac{1}{1-\alpha}\left(\frac{\bar{\beta}}{\bar{\beta}_2}\right)^{\frac{1}{1-\alpha}} - \frac{1-2\rho}{\rho}\left[1 - \left(\frac{\bar{\beta}}{\bar{\beta}_2}\right)^{\frac{1}{1-\alpha}}\right]\right\} > 0.
\end{aligned}$$

The first inequality holds using $\frac{2\rho-1}{1-(1-\rho)\rho\mu} > -1$. The third inequality clearly follows if $\rho \geq \frac{1}{2}$; if instead $\rho < \frac{1}{2}$, we have

$$\begin{aligned}
\frac{1}{1-\alpha}\left(\frac{\bar{\beta}}{\bar{\beta}_2}\right)^{\frac{1}{1-\alpha}} - \frac{1-2\rho}{\rho}\left[1 - \left(\frac{\bar{\beta}}{\bar{\beta}_2}\right)^{\frac{1}{1-\alpha}}\right] &> \frac{1}{1-\alpha}\left(\frac{\bar{\beta}_1}{\bar{\beta}_2}\right)^{\frac{1}{1-\alpha}} - \frac{1-2\rho}{\rho}\left[1 - \left(\frac{\bar{\beta}_1}{\bar{\beta}_2}\right)^{\frac{1}{1-\alpha}}\right] \\
&= \frac{1}{1-\alpha}[1-(1-\rho)\mu(1-\rho-\alpha+2\alpha\rho)] \geq 0,
\end{aligned}$$

where the last inequality uses $1-\alpha-\rho+2\alpha\rho \leq 1$ for $\alpha, \rho \in [0,1]$. Now observe that

$$f_a = \begin{cases} \frac{1}{2}\bar{\beta}[1-(1-\rho)\rho\mu]^{\alpha}, & \bar{\beta} < \bar{\beta}_1, \\ \frac{1}{2}\bar{\beta}^{\frac{1}{1-\alpha}}\left[\frac{\alpha(1-\rho)}{4}\right]^{\frac{\alpha}{1-\alpha}}, & \bar{\beta} \in \left[\bar{\beta}_1, \bar{\beta}_2\right], \\ \frac{1}{2}\bar{\beta}, & \bar{\beta} > \bar{\beta}_2, \end{cases}$$

from which clearly $\frac{\partial f_a}{\partial \rho} \lesseqgtr 0$ for $\bar{\beta} < \bar{\beta}_1$, $\frac{\partial f_a}{\partial \rho} < 0$ for $\bar{\beta} \in \left[\bar{\beta}_1, \bar{\beta}_2\right]$, and $\frac{\partial f_a}{\partial \rho} = 0$ for $\bar{\beta} > \bar{\beta}_2$. From $f_c = \sigma\rho\mu - u_0$ and $\frac{\partial \sigma}{\partial \rho} \geq 0$ we have $\frac{\partial f_c}{\partial \rho} = \mu\left(\frac{\partial \sigma}{\partial \rho}\rho + \sigma\right) > 0$. This completes the proof of statement (ii).

Consider statement (iii). First note that $\frac{\partial \bar{\beta}_1}{\partial \mu} < 0$ and $\frac{\partial \bar{\beta}_2}{\partial \mu} = 0$. This and $\left.\frac{\partial \sigma}{\partial \mu}\right|_{\bar{\beta} \in (\bar{\beta}_1, \bar{\beta}_2)} < 0$ imply $\frac{\partial \sigma}{\partial \mu} \leq 0$. From direct inspection of f_a above, we have $\frac{\partial f_a}{\partial \mu} < 0$ for $\bar{\beta} < \bar{\beta}_1$ and $\frac{\partial f_a}{\partial \mu} = 0$ otherwise. Finally, from $f_c = \sigma\rho\mu - u_0$ compute $\frac{\partial f_c}{\partial \mu} = \rho\left(\frac{\partial \sigma}{\partial \mu}\mu + \sigma\right)$. Notice $\frac{\partial \sigma}{\partial \mu}\mu + \sigma > 0$ when $\bar{\beta} < \bar{\beta}_1$ and $\frac{\partial \sigma}{\partial \mu}\mu + \sigma = 0$ when $\bar{\beta} > \bar{\beta}_2$. When $\bar{\beta} \in \left[\bar{\beta}_1, \bar{\beta}_2\right]$, we have

$$f_c = \frac{1}{(1-\rho)}\left[1 - \left(\frac{\bar{\beta}\alpha(1-\rho)}{4}\right)^{\frac{1}{1-\alpha}}\right] - u_0,$$

which does not depend on μ and implies $\frac{\partial \sigma}{\partial \mu}\mu + \sigma = 0$. To summarize, $\frac{\partial f_c}{\partial \mu} > 0$ when $\bar{\beta} < \bar{\beta}_1$ and $\frac{\partial f_c}{\partial \mu} = 0$ otherwise. This completes the proof of statement (iii) and Proposition 1. ■

Proof of Proposition 2: Consumer recommendation efficiency is decreasing in $\bar{\beta}$ since σ decreases in $\bar{\beta}$ from Proposition 1(i). An S-signalling consumer's payoff is $u_S - f_c = u_0 + u_S - \sigma\rho\mu$, and a D-signalling consumer's payoff is $\sigma\rho\mu + (1-\sigma)u_S - f_c = u_0 + (1-\sigma)u_S$. Both consumer types' payoffs are decreasing in σ and thus increasing in $\bar{\beta}$. It is immediate that platform profit increases in $\bar{\beta}$ because such a change increases n_a holding choice variables constant, hence optimal platform profit must increase. ■

Proof of Propostion 3: Let Π^* denote the platform's optimal profit, which from Proposition 1, is given by

$$\Pi^* = \begin{cases} \rho\mu - u_0 + \frac{1}{4}\bar{\beta}[1-(1-\rho)\rho\mu]^{\alpha}, & \bar{\beta} < \bar{\beta}_1, \\ \frac{1}{1-\rho}\left[1 - \left(\frac{\bar{\beta}\alpha(1-\rho)}{4}\right)^{\frac{1}{1-\alpha}}\right] - u_0 + \frac{1}{4}\bar{\beta}\left(\frac{\bar{\beta}\alpha(1-\rho)}{4}\right)^{\frac{\alpha}{1-\alpha}}, & \bar{\beta} \in \left[\bar{\beta}_1, \bar{\beta}_2\right], \\ -u_0 + \frac{1}{4}\bar{\beta}, & \bar{\beta} > \bar{\beta}_2. \end{cases}$$

From the proof of Proposition 1, we have $\frac{\partial \bar{\beta}_1}{\partial \rho}, \frac{\partial \bar{\beta}_2}{\partial \rho} > 0$. If $\bar{\beta} > \bar{\beta}_2$ before and after an increase in ρ, then $\frac{\partial \Pi^*}{\partial \rho} = 0$. If $\bar{\beta} > \bar{\beta}_2$ before an increase in ρ and $\bar{\beta} \in \left[\bar{\beta}_1, \bar{\beta}_2\right]$ after an increase in ρ, then $\frac{\partial \Pi^*}{\partial \rho} \geq 0$. This is because the platform optimally shifts from $\sigma = 0$ to $\sigma \in (0, 1)$, even though $\sigma = 0$ is still feasible with an unchanged profit. If $\bar{\beta} \in \left[\bar{\beta}_1, \bar{\beta}_2\right]$ before and after an increase in ρ, then compute

$$\frac{d\Pi^*}{d\rho} = \frac{1}{(1-\rho)^2}\left[1 - \left(\frac{\bar{\beta}\alpha(1-\rho)}{4}\right)^{\frac{1}{1-\alpha}}\right] \geq 0.$$

Finally, if $\bar{\beta} < \bar{\beta}_1$ after an increase in ρ, then again $\frac{\partial \Pi^*}{\partial \rho} \geq 0$. This is because Π^* is continuous in $\bar{\beta}$ and

$$\left.\frac{\partial \Pi^*}{\partial \rho}\right|_{\bar{\beta} \leq \bar{\beta}_1} = \mu\left(1 + \frac{\bar{\beta}}{\bar{\beta}_1}\,\frac{2\rho - 1}{1-\rho}\right) \geq 0,$$

using $\frac{2\rho-1}{1-\rho} \geq -1$.

Now consider an increase in μ. Recall that $\frac{\partial \bar{\beta}_1}{\partial \mu} < 0$ and $\frac{\partial \bar{\beta}_2}{\partial \mu} = 0$. If $\bar{\beta} < \bar{\beta}_1$, then directly compute $\frac{d\Pi^*}{d\mu} = \rho\left(1 - \frac{\bar{\beta}}{\bar{\beta}_1}\right) > 0$. Otherwise it is clear from inspection of Π^* that $\frac{d\Pi^*}{d\mu} = 0$ unless $\bar{\beta} \in \left[\bar{\beta}_1, \bar{\beta}_2\right]$ before and $\bar{\beta} < \bar{\beta}_1$ after an increase in μ, in which case $\frac{d\Pi^*}{d\mu} > 0$.

Consider statement (i). When $\sigma, \bar{\beta} > 0$, at a fixed recommendation quality $\rho\mu$, the platform strictly benefits from increased advertiser demand without any change to consumer demand from trading off μ for ρ when keeping its recommendation strategy and fees unchanged. If either σ or $\bar{\beta}$ is zero, then such a change has no effect on platform profit.

To see statement (ii), notice that for sufficiently high $\bar{\beta} \in \left[\bar{\beta}_1, \bar{\beta}_2\right]$, $\frac{\partial \Pi^*}{\partial \mu} = 0$, while $\frac{\partial \Pi^*}{\partial \rho} \geq 0$ and strictly so within this interval. ■