2021

Essays on the design of innovative subscription-based business models

https://hdl.handle.net/2144/43127

BOSTON UNIVERSITY

QUESTROM SCHOOL OF BUSINESS

Dissertation

ESSAYS ON THE DESIGN OF INNOVATIVE SUBSCRIPTION-BASED BUSINESS MODELS

by

ELNAZ KARIMI

B.Sc., Sharif University of Technology, 2011
M.Sc., Concordia University, 2013
Ph.D., Polytechnique Montreal, 2018

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

2021

Approved by

First Reader

Pnina Feldman, PhD
Associate Professor of Operations and Technology Management

Second Reader

Anita L. Tucker, Ph.D.
Everett W. Lord Distinguished Faculty Scholar
Professor of Operations and Technology Management

Third Reader

Ella Segev, Ph.D.
Senior Lecturer
Ben Gurion University of the Negev

Acknowledgments

This thesis is the product of endless nights, stressful hours, and countless work iterations. It would not have been possible without the dedicated support of many individuals.

First, my deepest gratitude goes to my advisor Prof. Pnina Feldman. I visited Pnina's office with a model in the first week I met her. She listened carefully, then heavily criticized the model, and finally told me that I "need to think about the micro-foundations of decision making." Looking back at that day and my first models, I realize how far I have come, and how much of it I owe to Pnina's selfless advice and patience. Academically, Pnina's immense knowledge of Operations Management, combined with her experience as an economic modeler and researcher provided the fundamentals I needed to grow as a researcher. On a personal level, Pnina's support and words always comforted me at the times of stress, and her guidance helped me navigate my PhD journey.

Next, I would like to thank Prof. Ella Segev for the many hours she spent to read my models, explain how economists model, and provide feedback on my work. I will always appreciate how much you taught me.

I would like to thank Prof. Anita Tucker, our PhD liaison, for her support. I cannot thank Anita enough for her optimistic yet objective way of looking at problems, and her selfless advice. Talking to Anita always shed light and hope, and gave me the strength to take yet another step forward.

I also want to thank the faculty of Questrom - Tim Simcoe, Janelle Heineke, Nitin Joglekar, Joe Restuccia and Marcus Bellamy - thank you for all your support and for helping me grow into the person I am now.

To my fellow PhD students - Brad Lee, Emre Guzelsu, Alf Wang, Temidayo Adepoju, Minje Park, and Ankita Shirahatti - thank thank you for being such amazing friends and making our time together so special and fulfilling .

I wish to thank my family for always being there. Despite the thousands of miles between us, I felt your presence and support every single moment. To my mom, Fatemeh, and my dad, Alireza, there are no words that can express my gratitude and love for you. I would not be the person I am today without your unconditional love and support. To my dearest sister, Behnaz, thank you for being there for me every step of the way and for believing in me even when I didn't. To my brother, Mehdi, thank you for making me laugh and reminding me once in a while there is life outside of my doctoral program. To my parents-in-law, Mohanna and Mohammadreza, and my brother-in-law, Aarash thank you for believing in me and cheering me on all steps of the way. I love you all so much.

Finally, to my husband and best friend, Arshya: It has been a blessing to have you by my side every step of the way. Your kindest words when I felt in trouble gave me hope, and your encouraging words when I wanted to settle for anything less than "perfect" reminded me to never underestimate myself and never set a limit for what I can achieve. You have always been my harshest critic, yet biggest fan; and I thank you for the endless love you have given me since I have known you.

ESSAYS ON THE DESIGN OF INNOVATIVE SUBSCRIPTION-BASED BUSINESS MODELS

ELNAZ KARIMI

Boston University, Questrom School of Business, 2021

Major Professor: Pnina Feldman, PhD
Associate Professor of Operations and Technology Management

ABSTRACT

This dissertation consists of three essays that study operations management problems of emerging subscription-based business models. The first essay, coauthored with Pnina Feldman and Ella Segev, investigates the personalized subscription box business model where a firm selects and delivers products that are curated to match each customer's tastes and preferences. Specifically, this essay investigates how firms may optimally learn customer tastes, and personalize their offerings in order to maximize their revenue. Doing so, we investigate the optimal pricing and return policies of a firm operating with the personalized subscription box business model.

The second essay in this dissertation studies the coordination problem of a two-sided media streaming platform that aggregates first and third-party content in a bundle and offer the bundle to users for a subscription fee. Media streaming platforms face a trade-off between maximizing their user base by attracting premium content providers and the costs due to the outside options of the providers. We study the effectiveness of platform's strategic investment in first-party content to attract premium providers under conventional revenue-allocation mechanisms.

The third essay in this dissertation studies two-sided media streaming platforms that offer both first- and third-party content, but strategically steer users to the first-party content, which is most profitable to the platform. Platforms do so by manipulating their recommendation systems. In this case the platform may choose to recommend a content that is different from a user's optimal preferences. The purpose of this paper is two-fold. First, we explore the effect of platform's first-party content bias on users' search behavior. Second, we study the effect of users' search cost and the third-party provider's royalty fee on the design of recommendation systems.

Contents

List of Figures

Chapter 1

Introduction

1 Overview

Because of consumers growing desire for a convenient and personalized experience, many incumbents and new entrants in e-commerce and digital goods industry trend towards a bundling and subscription revenue model. This dissertation focuses on two such business models: *personalized (curated) subscription box services* and *media streaming services*. The goal of this dissertation is to study the consumer behavior and the strategic decisions of innovative subscription-based platforms, and examines how to enhance efficiency and revenues of these business models by providing actionable operational policy recommendations. To study these models, we employ a variety of analytical techniques from operations management and economics, including stylized economic modeling and game theory.

The second chapter studies the *personalized (curated) subscription box* business model. In recent years, led by start-ups such as Blue Apron, Trunk Club and Stitch Fix, personalized subscription services have become increasingly popular amongst online shoppers. In fact, a survey conducted by McKinsey & Co. shows that 15 percent of online shoppers have signed up for one or more subscriptions to receive products on a recurring basis[1]. In a subscription box business model the firm selects and delivers products that are curated to match each customer's tastes and interests.

---

[1] https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/thinking-inside-the-subscription-box-new-research-on-ecommerce-consumers

Whereas in a traditional shopping experience a customer chooses the product(s) she wishes to purchase from an assortment of options, in a personalized subscription box service, she leaves the selection process to the firm. Therefore, such purchases can be perceived as riskier by customers. Many subscription services allow customers to return unwanted items to mitigate this risk. An important mechanism that echoes the uniqueness of subscription services is learning and personalization. Subscription services have given an unprecedented ability to firms to learn their customers' tastes and tailor their services accordingly.

The third chapter studies the coordination problem of media streaming platforms. Media streaming services, such as Netflix, Hulu and Amazon Prime Video, aggregate a large number of content produced by third-party providers and offer unlimited access to the bundle to users for a subscription fee. These platforms act as two-sided content aggregators and benefit from strong indirect network effects. Namely, the more content on the platform, the more users value and join the service and, in turn, content providers have higher incentives to join a platform in which they can meet a wider audience. However, not all content generate the same externalities. The presence of popular tv-shows (e.g. The office, Friends) or classic movies (e.g. The Godfather, Singin' in the Rain) are generally valued more by users than their niche counterparts. Therefore, providers that own these "popular content" demand to be compensated for the externalities they create. In short, for these platforms, managing the supply side of the market - i.e. attracting the "popular content" producers - can be as crucial as managing the demand side. However, little is known about the platform's coordination challenge and the consequences of heterogeneity on the supply side on the ability of the revenue allocation rules to coordinate the market. Furthermore, the higher price of popular content has incentivized the streaming platforms to produce their own content. Though creating first-party content has proven to be more expensive and riddled with uncertainty than licensing third-party content, many formerly agrregator streaming platforms including Netflix, Amazon Prime and Hulu have invested significantly to produce first-party content. Therefore, streaming platforms have transformed to a hybrid business model, operating as a third-party content agrregator and a first-party content producer, which further complicates the coordination problem of these two-sided platforms.

Finally, in the fourth chapter we study the design of personalized recommendation services of media streaming platforms. Two-sided media streaming platforms that offer both first- and third-party content may be inclined to strategically steer users to the first-party content, which is most profitable to the platform. Platforms do so by manipulating their recommendation systems such that the platform recommends a content different from a user's optimal preferences. Often the streaming services such as Netflix carry a large number of content in their library and without recommendation services, the search may be very costly to users. Therefore, a large user segment of such platforms rely on the recommendation of engines to choose what content to consume. For example, according to reports published by Netflix, 75% – 80% of viewer activity on Netflix is influenced by its recommendation algorithm[2]. However, users are aware of the platform's recommendation bias, and consider the accuracy of recommendations. Therefore, bias may influence users' optimal search policy and result in lower willingness to pay.

---

[2] How Netflix Uses Analytics To Select Movies, Create Content, and Make Multimillion Dollar Decisions [Internet]. Available from: https://blog.kissmetrics.com/how-netflix-uses-analytics/.

2 Summary of Results

2.1 Chapter 2: The Subscription Box Business Model: Learning Tastes and Product Returns

In this paper, we investigate how a subscription service firm should utilize its operational levers to learn a strategic customer's taste and personalize its services to her preferences. We show the firm faces a trade-off between exploration (learning preferences) by sending risky products versus exploitation (earning revenue) by sending safe products. We find that depending on the hassle cost of receiving an unwanted product, it can be optimal for a firm to experiment with its content. Furthermore, we show that when facing forward-looking customers, a fixed pricing strategy performs better than a dynamic pricing strategy. The fixed-pricing strategy not only yields higher expected revenue and allows the firm to experiment with higher hassle costs of return, but also, maximizes customer acquisition. Finally we show that it is optimal for the firm to take away all the hassle cost of return, thus making returns hassle-free for the customer.

2.2 Chapter 3: Coordination of Streaming Platforms in the Presence of First-Party Content

In this paper, first, we explore the performance of well-established revenue allocation rules (*revenue-sharing* and *licensing*) in coordinating the two-sided platform when there exists heterogeneity on the supply-side. Second, we suggest a rationale to the platform's choice to invest in the risky production of first-party content and the impact on its ability to attract popular content providers. We show when content providers are heterogeneous in their popularity amongst users, revenue-sharing does not coordinate the market. The reason is that the platform is inclined to manipulate the consumption rate of the popular content to decrease the payment to the popular providers. However, we show that if the platform commits to charge a subscription fee, the revenue-sharing contract can coordinate the market. Furthermore, we find licensing is able to coordinate the market, because the platform pre-commits to a payment to the providers before setting the subscription fee. Finally, we find the platform invests more aggressively in the production of its first-party content, to be able to attract the popular provider to join the platform.

2.3 Chapter 4: First-Party Bias in Media Streaming Platforms

This paper studies the design of the recommendation system of an streaming platform that has private information about users valuations and provides users with personalized recommendations. We show depending on the intensity of platform's bias to recommend first-party content, users update their belief about their valuation for alternative content and may choose to continue sampling to find the content that best matches their tastes. We show a biased recommendation system allows the platform to increase its revenue, albeit at the cost of lower social welfare. We also find that as the search cost increases the platform is able to steer users toward first-party content more aggressively. Furthermore, we find that when the search cost is high, the third-party provider is better off with an intermediate royalty rate.

Chapter 2

The Subscription Box Business Model: Learning Tastes and Product Returns

1 Introduction

Recent years have witnessed a substantial increase in the popularity of e-commerce subscription services. In a subscription box business model, the firm selects and delivers products curated to match each customer's tastes and interests, in regular time intervals. Subscription boxes accounted for approximately \$15B in U.S. sales in 2018, and have established a foothold in various categories such as beauty, fashion, food and home decor[1]. Notable examples include Birchbox, Trunk Club and Stitchfix.

Whereas in a traditional shopping experience a customer chooses the product(s) she wishes to purchase from an assortment of options, in a personalized subscription box service, she leaves the selection process to the firm. A key challenge for firms is to learn individual customer tastes and personalize their subscription box contents, accordingly. Personalization requires deep insight into customer behavior, which can only be achieved through educated guesses, and repeated interactions with the customer (Aydin and Ziya, 2009). However, the obvious cost of learning customer preferences through trial and error is the risk of sending too many unwanted products, and consequently, losing the customer. In fact, dissatisfaction with assortment is one of the main reasons customers cancel their subscription service (Chen et al., 2018).

---

[1] See https://www.forbes.com/sites/andriacheng/2019/12/20/the-crowded-subscription-box-industry-still-holds-growth-promise/?sh=2a445bbb21ae.

To mitigate this risk, most firms offer customer-friendly features such as free service cancellation and free returns in the event that a customer receives an unwanted product[2]. This paper investigates how firms may optimally learn customer tastes, and personalize their offerings in order to maximize revenue. Doing so, we investigate the optimal pricing and return policies of a firm operating with the personalized subscription business model.

In a subscription setting, the firm presents a customer a small subset of their assortment options, observes the customer's purchasing choice, and reacts accordingly. In determining the small subset, the firm faces a trade-off. On one hand, offering a variety of diversified assortment with uncertain values and observing the customer's responses facilitates learning the customer's specific preferences through experimentation. The main advantage of learning is that the firm may adjust its pricing and assortment choices dynamically to extract more surplus from the customer over time. On the other hand, extensive (failed) experimentation may result in frequent returns or customer abandonment, and therefore, compromise the firm's revenue. Hence, the firm must also consider the possibility of presenting the customer with a less risky option, for which the firm is guaranteed to receive *some* revenue. In short, a subscription firm has two personalization levers: 1) selecting an assortment to present to a customer; 2) pricing.

In this paper, we investigate how learning customer tastes through experimentation affects a firm's content and price personalization decisions. Specifically, we address the following research questions. (1) When should firms attempt to learn customer tastes through experimentation? (2) How does a customer's strategic behavior affect a firm's actions? (3) How do a firm's pricing and return policies affect its revenue and ability to learn? (4) How does experimentation affect consumer surplus,

---

[2] While a few clothing subscription services such as Trunk Club allow customers to preview the box's assortment prior to delivery, most firms send the products sight unseen.

and social welfare?

To address these questions, we examine the interaction between a subscription service firm and a single forward-looking strategic customer using a two-period stochastic game-theoretic model. The firm is in an abundant possession of products from two distinct categories: “safe” and “risky”. The customer’s preference for products from the safe category is known to both the firm and the customer. However, the customer’s preference for products from the risky category is private information to the customer, and unknown to the firm. Importantly, we assume that the customer values a risky product that she likes higher than a safe product. Our model setting is in the spirit of the seminal two-armed bandit model (e.g., see DeGroot (2005), section 14.5). Albeit, we tailor the model to account for information asymmetry between the firm and the customer. Whereas in a typical two-armed bandit model, the principal directly observes the payoff from pulling an “arm”, in our model, the strategic customer is the party that observes her value and signals this information to the firm through her actions. The firm makes content assortment and pricing decisions. Specifically, in each period, the firm chooses to send a product from either the safe or risky category. With regards to pricing, the firm may employ one of two policies: the firm may either: (a) announce the prices for each category at the beginning of each period (dynamic pricing); or (b) set fixed prices for each category at the beginning of the subscription (fixed pricing). An important aspect of personalized subscription services, which is captured in our model, is the hassle cost of returns. The costs of making a return can be personal hassle costs, or penalties such as return shipping or restocking fees imposed by the firm. We assume this cost is exogenous. We model the customer’s purchasing choice such that upon realizing her value for a product she receives, she may choose to keep it and pay its price, or return the product and incur the hassle cost of return.

Our equilibrium analysis of this model yields several important insights and managerial implications. First, we find that under either pricing policy, when the hassle cost of return is below a certain threshold, the firm engages in experimentation to learn the customer's taste. Specifically, the firm sends a risky product to the customer in period 1 to partially learn her preference for contents from the risky category. Following the experimentation phase in period 1, the firm employs behavior-based content personalization in period 2. Namely, in response to observing that the customer keeps the product in period 1, the firm sends a risky product in period 2. As the hassle cost of return increases, the cost of learning outweighs its benefit, and the firm always sends safe products, and forgoes experimentation. The reason is that if the hassle cost of return is high, the risky products must be priced relatively low to convince the customer to join the service. Consequently, the firm's expected revenue from offering risky products is low, rendering experimentation futile.

Second, we find fundamental differences between personalization under dynamic pricing and fixed pricing. We find that when the firm experiments under a dynamic pricing policy, it exploits both levers of content and price personalization. In period 1, to ensure that a transaction is made (i.e. the customer joins the service), the firm offers the risky products at a lower price. After period 1's experiment with sending a risky product, if the firm believes that the customer is likely to have a favorable taste for the risky category, then the firm increases the price (i.e. price personalization) and sends another risky product in period 2 (i.e. content personalization). However, if the customer returns a risky product, the firm switches to the safe category. This result is in line with the literature on behavior-based discrimination in repeatedly-purchased products (Murthi and Sarkar, 2003). Under fixed pricing, however, the firm's only lever of personalization is content. We find that under fixed pricing, if the firm chooses to experiment, it discounts the price of the safe category, but increases the price of the risky category. This pricing strategy enables the firm to convince the customer to join the service, regardless of her preference for risky products. Specifically, a customer with a distaste for risky products – who would not otherwise join the service under dynamic pricing – participates in the experimentation phase, *anyway*, anticipating that once the firm learns her true preferences, she will receive safe products, but with a discount. Surprisingly, we find that although dynamic pricing offers two levers of personalization, the firm is better off with a fixed pricing policy. This is because under fixed pricing, the customer is guaranteed to join the service and also, the fixed pricing strategy allows the firm to engage in experimentation for greater values of the hassle cost of return. Finally, we find that fixed pricing achieves higher total social welfare since it is not only beneficial for the firm, but is also preferred by the customer.

Finally, we find that it is optimal for the firm to take away all the hassle cost of return, thus making returns hassle-free for the customer. In practice, most personalized subscription services offer full refunds (100% money-back guarantee) for returns. Additionally, firms may also adopt different policies that determine how convenient it is for the customers to return products. Such policies include, but are not limited to, offering generous time limits, free shipping for returns, free at-home pickups, etc., which facilitate customer returns. The benefit of hassle-free returns is twofold. First, the firm can better extract value from the customer by charging a higher price for risky products. Second, the firm ensures revenue by guaranteeing that the customer joins, and a transaction is made. Interestingly, although total social welfare increases with hassle-free return, we find that customers are worse-off.

2 Literature Review

In this paper, we focus on the design of personalized subscription services, which are characterized by a firm's attempt to learn a strategic customer's taste. Specifically, we explore pricing policies and assortment selection, as the firm's levers for personalization. Furthermore, we investigate the firm's optimal return policies when facing strategic customers. As such, our paper lies at the intersection of three streams of research: demand (preference) learning, personalization, and strategic customer behavior.

The literature on demand learning through assortment planning and pricing policies is relevant to our work. The studies on assortment planning mostly consider a homogeneous population of customers with unknown demand. For example, Caro and Gallien (2007) uses Bayesian inference in a multi-armed bandit context to dynamically learn the demand distributions in a multi-product setting. Rusmevichientong et al. (2010) studies the assortment planning problem with capacity constraints using a Multinomial Logit (MNL) model and presents an "explore first and exploit later" approach. (See Strauss et al. (2018) for a thorough review of assortment optimization with demand learning.) Closely related to our work is Bernstein et al. (2020) that uses a stylized model to show that subscription services can employ a "dynamic content variation" strategy to maximize their demand over a predetermined planning horizon. They find that when customer heterogeneity is low (high) the firm must increase (decrease) the quality of its content to ensure customer retention (acquisition). There exists a rich literature on demand learning via price experimentation that utilizes the multi-armed bandit algorithm to study the trade-off between exploration (demand learning) and exploitation (revenue earning). Aviv and Vulcano (2012) provides a survey of works in this area. In a more recent work, Hu et al. (2019) studies the joint demand learning and pricing problem in a model where consumers can return the product. Our paper contributes to this literature by introducing a novel learning model that can be applied to firms that have repeated one-to-one interactions with their customers.

A stream of literature in marketing investigates how firms can exploit customers' personal information for the purpose of personalized pricing (Candogan et al., 2012), quality discrimination (Li, 2021) and targeted advertising (Shen and Miguel Villas-Boas, 2018). Prompted by the growing availability of data in online retailing, personalized assortment planning and personalized pricing have become an active research area in revenue management. (See Mišić and Perakis (2020) for a comprehensive review.) The existing literature in this area typically assumes that customers are price-sensitive myopic shoppers, and designs algorithms that aim to maximize firms' expected revenue through personalized assortment and pricing. For instance, Bernstein et al. (2015) studies a firm with limited inventory that customizes the subset of products it offers to its customers, based on their preferences. It derives the structural properties of the optimal policy and proposes heuristics for assortment customization. Cheung and Simchi-Levi (2017) examines an online personalized assortment optimization problem with unknown purchase probabilities and dynamic pricing. It proposes an algorithm to estimate latent parameters of an MNL choice model, and evaluates the performance of the algorithm using Bayesian regret. Bernstein et al. (2019) integrates assortment decisions with online (i.e. real-time) demand learning using a clustering-based policy. This paper assumes that there exists an unknown number of clusters of customer profiles, and uses the Dirichlet process mixture model to estimate the cluster distributions. In a similar approach, Ban and Keskin (2020) uses customer data to develop an algorithm for dynamic pricing of a single product at an individual customer level. More closely to our work, Aydin and Ziya (2009) considers the case of a personalized dynamic pricing scenario, where customers belong to either a high or low reservation price group. The firm learns its customers' preferences through signals they send, which indicate how likely they are to have a high reservation price. Our work is distinct from the aforementioned studies because we study assortment and price personalization in the context of a subscription service, where, essentially, the firm is the party that governs the customer's purchasing decision through individual customer-level content choices. Moreover, we investigate the firm's operational levers in the presence of strategic consumer behavior.

There exists a rich literature related to strategic consumer behavior in operations and revenue management. The impact of strategic consumer behavior on a firm's operations including inventory-related strategies (Su and Zhang, 2008), product design (Feldman et al., 2019), and pricing strategies such as dynamic pricing (Cachon and Swinney, 2009), fixed pricing (Su and Zhang, 2008), and price commitment (Aviv and Pazgal, 2008), have been extensively studied. Most related to our work are papers that investigate a firm's optimal return and pricing policies. Su (2010) studies a dynamic pricing problem for repeatedly-purchased products. The paper shows that a fixed pricing strategy can deter the stockpiling behavior of strategic customers, and that dynamic pricing should only be used when frequent buyers have a high willingness to pay. Papanastasiou and Savva (2017) considers a setting in which strategic consumers are allowed to delay their purchases to learn about a product's quality through online reviews. It finds that in the presence of social learning, a responsive (i.e. dynamic) pricing strategy is preferred over a price-commitment pricing strategy to deter strategic waiting. Altug and Aydinliyim (2016) study a setting in which strategic consumers are uncertain about their valuations early in the season, and postpone their purchase to a later period when their uncertainty is resolved, and price is discounted. It shows that allowing returns with a generous refund could deter strategic waiting behavior, and induce more sales at a higher price at the beginning of the season. We contribute to this literature by examining the implications of strategic consumer behavior in a subscription service. We show that in such settings, fixed pricing with product returns can be a more effective mechanism for learning customer preferences than dynamic pricing. This is because a forward-looking and rational customer is willing to allow the firm to learn her taste, in exchange for a future discount on products that better match her taste.

3 Analysis

3.1 Model Description

We study a firm that offers a subscription service to a customer over two periods. The firm offers two types of product categories, each with zero marginal and fixed costs. The first type is an established product category, which we refer to as the safe category. A safe product has known value to the customer, which practically implies that it is risk-free to both the firm and the customer. The second type is appealing to some – but not all – customers, which we refer to as the risky category. In other words, a risky product has uncertain value to the customer and thus, imposes a risk to both parties. At the beginning of each period, the firm chooses from which category to send a product, and at what price. The prices are denoted as (p_1^s, p_1^r) and (p_2^s, p_2^r) for the safe and risky products in period 1 and period 2, respectively. We assume that prices are set dynamically. In other words, the firm sets the price in each period to maximize profit-to-go.

At the beginning of each period, the customer must choose whether she wishes to join the service or not; joining (not joining) in the second period implies resuming (leaving) the service. We assume that the customer can choose to join the service at no cost and is free to leave the service at any point. At the beginning of each period, the customer receives a product from one of the two categories. She derives value v from the product, and opts to keep it and pay price p, or return it and incur a hassle cost h $(0 \leq h \leq \frac{1}{2})$. The hassle cost of return captures the fact that returns can be inconvenient to a customer for reasons such as traveling to return the product, the waiting time to be refunded, or due to shipping costs. If the customer decides to keep a product of value v at price p, her immediate utility is $u(v) = v - p$. If she decides to return the product, her immediate utility is $u = -h$. We further make the assumption that the customer keeps the product if she is indifferent between keeping/returning, and similarly, joins the service if she is indifferent between joining and not joining. Moreover, the customer is risk neutral with a zero outside option; hence, she joins the service if her expected utility from joining is non-negative. The customer has a fixed value of $v = \frac{1}{2}$ for all safe products, but is uncertain about her value for risky products. Specifically, she may like a risky product and realize a value of $v = 1$ with a probability of δ, or dislike the product and realize a value of $v = 0$, with a probability of $1 - \delta$. In addition, we assume the customer's valuations for risky products are i.i.d across periods. While the customer has private knowledge of her δ (i.e., she knows how often she likes risky products), the provider has no information. Furthermore, we make the assumption that the customer is strategic with respect to her decision to join in the first-period or resume in the second period (i.e. she is rational and considers the future when making her decision to join/resume the service). However, to simplify the analysis, we assume that the customer is myopic with respect to her decision to keep/return a product she receives. Therefore, she does not consider period 2 when making her decision to keep/return a product in period 1; consequently, she keeps a product in period 1 if and only if her immediate utility from keeping the product is greater than that of returning the product. Finally, we assume that the firm begins the service with a prior belief that δ is uniformly distributed on the interval $[0, 1]$.

The timeline of the game is as follows: at the beginning of period 1, the firm announces the prices for the safe and risky products. The customer observes (p_1^s, p_1^r) and decides whether or not to join the service. If the customer chooses to join, the firm must decide whether to send a product from the safe or risky category. Once the customer receives the product, she realizes her value and must decide to keep or return it. Upon observing the customer's action, the firm updates its belief about the customer's δ. With the updated belief, the firm announces the new prices. The customer observes the prices and may choose to leave the service, at which point the game ends. If the customer chooses to stay, the game proceeds into the second period, and once again, the firm must choose from which category to send a product.

3.2 Dynamic Pricing Equilibrium

We analyze the two-period stochastic game between the firm and the customer using backward induction. We seek to find a pure-strategy Bayesian perfect equilibrium.

Special Case: Hassle-free Return

Before solving for the general equilibrium, it is useful to start with the firm's product and pricing choice in a special case where returns are hassle-free for the customer (i.e. $h = 0$). Lemma 2.1 characterizes the perfect Bayesian equilibrium, when $h = 0$. We find when returns are hassle-free for the customer, the firm is always better off by sending a risky product in period 1 and learning the customer's taste. We present the proof in the main text because it is simple and motivates our approach to the general case.

Lemma 2.1. *When returns are hassle-free for the customer, the firm's unique optimal pricing scheme is such that:* $p_1^{r*} = p_2^{r*} = 1$ *and* $p_1^{s*} = p_2^{s*} = \frac{1}{2}$. *Moreover, the firm's unique optimal product choice is such that:*

- *(i) In period 1, the firm sends a risky product.*
- *(ii) In period 2, the firm sends a risky product if it observes "keep" in period 1, and sends a safe product if it observes "return" in period 1.*

The customer joins the service in period 1, and resumes the service in period 2 regardless of her δ. *In both periods, she keeps a product she likes, and returns it, otherwise. The firm's optimal expected revenue from the service is* $\pi^* = \frac{13}{12}$.

Proof. Following backward induction, we first examine the customer's decision in period 2, when receiving a risky or safe product. The customer keeps the product if her utility is non-negative and returns it at no cost, otherwise. Therefore, the customer keeps a risky product in period 2 if $v - p_2^r \geq 0$, and her expected utility from receiving a risky product is $U_2(r) = \delta\,(1 - p_2^r)$. Recall that δ is the probability that the customer will like a risky product. Similarly, the customer's utility from receiving a safe product is $U_2(s) = \frac{1}{2} - p_2^s$, and she keeps the product if $p_2^s \leq \frac{1}{2}$. Also, the customer resumes her service as long as her utility is non-negative, or $p_2^r \leq 1$ and $p_2^s \leq \frac{1}{2}$.

The firm's optimal pricing strategy in period 2 is as follows: the firm prices a risky product at $p_2^{r*} = 1$, and a safe product at $p_2^{s*} = \frac{1}{2}$. Note that ensuring that the customer resumes the service (i.e. $p_2^r \leq 1$ and $p_2^s \leq \frac{1}{2}$), the firm's revenue is strictly increasing in price. Therefore, it is optimal for the firm to price the product it sends at the customer's reservation price. Let us assume that the firm begins period 2 with the posterior pdf $f(\cdot)$ with support $[0, 1]$. Then the firm sends a risky product if the revenue from sending a risky product priced at $p_2^{r*} = 1$, $\pi_2^r(p_2^{r*}) = \int_0^1 s f(s)\, ds$, is greater than the revenue from sending a safe product, $\pi_2^s = \frac{1}{2}$. Therefore, the firm sends a risky product if $\int_0^1 s f(s)\, ds > \frac{1}{2}$.

In period 1, similar to period 2, the customer keeps a product if her utility is non-negative and returns it, otherwise. Moreover, she joins the service if her expected utility from receiving a product from either categories is non-negative. Therefore, analogous to period 2, the firm prices a risky product at $p_1^{r*} = 1$, and a safe product at $p_1^{s*} = \frac{1}{2}$. Next, we consider the firm's product choice in period 1. If the firm sends a risky product in period 1, the customer keeps it if her realized value is 1 and returns it, otherwise. Therefore, according to Bayes rule, upon observing the customer's action to keep (Equation 2.1) or return (Equation 2.2) the product, the firm updates its belief as follows:

$$f(\delta \mid \text{keep}) = \frac{p(v = 1 \mid \delta)\, f(\delta)\, d\delta}{\int_0^1 \delta f(\delta)\, d\delta} = 2\delta \tag{2.1}$$

$$f(\delta \mid \text{return}) = \frac{p(v = 0 \mid \delta)\, f(\delta)\, d\delta}{\int_0^1 (1 - \delta) f(\delta)\, d\delta} = 2\,(1 - \delta) \tag{2.2}$$

The firm's revenue from sending a risky product in period 2 following a keep (return) is $\pi_2^r = \int_0^1 2s^2 ds = \frac{2}{3}$ $\left(\pi_2^r = \int_0^1 2(1-s) s ds = \frac{1}{3}\right)$. Therefore, the firm sends a risky product following a keep observation, and a safe product following a return observation. □

General Case

Second-period subgame: Following backward induction, we first look at the period 2 subgame. In period 2, if the customer receives a safe product, she keeps it if she gains a higher utility by keeping the product rather than returning it. Recall that the customer incurs a hassle cost of h should she choose to return the product. Therefore, the customer keeps a safe product in period 2 if $\frac{1}{2}-p_2^s \geq -h$, which implies $p_2^s \leq \frac{1}{2}+h$. In this case, the customer's utility from receiving a safe product in period 2 is $U_2(s) = \frac{1}{2} - p_2^s$. The utility function implies that the customer resumes the service if $p_2^s \leq \frac{1}{2}$. If the customer receives a risky product, she realizes her value upon observing it, and keeps the product if $v - p_2^r \geq -h$. Therefore, if $p_2^r \leq h$, it is too costly for the customer to return a product she does not like, and will keep the risky product regardless of her value. In this case, she resumes the service at the end of period 1 if her $\delta \geq p_2^r$ because this would imply that her utility from resuming the service and receiving a risky product in the second period is $U_2(r) = \delta - p_2^r \geq 0$. However, if $h < p_2^r \leq 1+h$, the customer returns the product she does not like. Hence, she will like the product, keep it and get a utility of $(1 - p_2^r)$ with probability δ, or dislike it and incur the hassle cost of return (h) with probability $1-\delta$. It follows that her expected utility from receiving a risky product in period 2 is $U_2(r) = \delta\,(1 - p_2^r) - (1-\delta)\,h$, and the customer resumes the service in period 1 if $U_2(r) \geq 0$, or $\delta \geq \frac{h}{1-p_2^r+h}$. (Note that if $p_2^r > 1$ the customer never resumes the service regardless of her δ, since the price of the product is greater than the maximum value for the product, $v = 1$, and resuming entails a negative expected utility.)

To continue the analysis of the second period subgame, we temporarily assume that if the customer has joined the service in period 1, it entails that she has a $\delta \geq \delta_1$, where $0 \leq \delta_1 \leq 1$. The validity of this assumption is proven in the next section. Intuitively however, upon observing the firm's prices (p_1^s, p_1^r) in period 1, if a customer with δ_1 who is indifferent between joining or not joining the service chooses to join, then it stands to reason that a customer with a $\delta > \delta_1$ would also join the service. Moreover, as a result of observing the customer's first-period actions, the firm begins period 2 with the posterior pdf $f(\cdot)$ with support $[\delta_1, 1]$. The firm must announce the prices and choose the product category to send to the customer.

Let π_2^s and π_2^r be the firm's revenue from sending a safe product and a risky product in period 2, respectively. Recall that the customer resumes the service at the beginning of period 2 if her utility is non-negative. This occurs if the customer is to receive a safe product (regardless of her δ), or if she is to receive a risky product and her $\delta \geq \frac{h}{1-p_2^r+h}$. For receiving a risky product in period 2, let us denote the customer's resuming threshold as $\delta_2 = \max\left(\frac{h}{1-p_2^r+h}, \delta_1\right)$. With this in mind, if the firm opts to send a safe product with a price $p_2^s \leq \frac{1}{2}$ (to ensure a non-negative utility for the customer), the customer will keep the safe product she receives and pays p_2^s. Therefore, the revenue from sending a safe product is:

$$\pi_2^s(p_2^s) = p_2^s$$

and the optimal price is $p_2^{s*} = \frac{1}{2}$. Alternatively, if the firm opts to send a risky product with $h < p_2^r \leq 1$, the customer resumes the service at the beginning of period 2 if her $\delta \geq \delta_2$. Hence, the firm's expected revenue is derived based on the joint probability that the customer resumes the service and keeps the product, and the price of the product. Therefore,

$$\pi_2^r(p_2^r) = p_2^r \int_{\delta_2}^{1} s f(s)\, ds$$

Let $p_2^{r*} = \arg\max \pi_2^r(p_2^r)$ be the firm's optimal price for risky products in period 2. The firm sends a safe product if $\pi_2^s(p_2^{s*}) \geq \pi_2^r(p_2^{r*})$. Lemma 2.2 establishes the firm's optimal product and price choices in period 2, and the customer's optimal response. Moreover, for any given $f(\cdot)$, the equilibrium in the second-period subgame is unique.

Lemma 2.2. *For any posterior pdf $f(\cdot)$, there exists a unique (p_2^{s*}, p_2^{r*}) that maximizes the firm's expected revenue from sending a risky or a safe product in the second period. Given (p_2^{s*}, p_2^{r*}), there exists a unique subgame-perfect equilibrium in the second period played by the firm and the customer as follows:*

(i) If $\pi_2^r(p_2^{r*}) > \pi_2^s(p_2^{s*})$, then (a) the firm announces (p_2^{s*}, p_2^{r*}); (b) the customer resumes the service if her $\delta \geq \delta_2(p_2^{r*})$; (c) the firm sends a risky product; (d) the customer keeps the product if her realized value is 1, and returns it otherwise.

(ii) If $\pi_2^r(p_2^{r*}) \leq \pi_2^s(p_2^{s*})$, then (a) the firm announces (p_2^{s*}, p_2^{r*}); (b) the customer resumes the service; (c) the firm sends a safe product; (d) the customer keeps the product.

All proofs are provided in Appendix A.

First-period subgame: We now bring the preceding analyses together and consider the implications of customer behavior for the firm's product choice and expected revenue.

If the firm sends a safe product in period 1, the customer always keeps it, and because of this, sending a safe product in period 1 never provides the firm with any information regarding the customer's taste δ; and in the absence of information about δ, it is optimal for the firm to send a safe product. Therefore, if the firm chooses to send a safe product in period 1, the firm sends another safe product in period 2. Under this subgame, $p_1^{s*} = p_2^{s*} = \frac{1}{2}$, and the firm's optimal expected revenue is $\pi_1^*(s) = 1$. The customer joins regardless of her δ, and the firm extracts all customer surplus.

If the firm chooses to send a risky product in period 1, the customer joins if her $\delta \geq \delta_1$. She keeps the risky product if her realized value is $v = 1$ and returns it, otherwise. Upon observing the customer's response, following the Bayesian update function described in equations 2.1 and 2.2, the firm updates its belief as follows:

$$f\left(\delta \mid \text{keep}\right) = \frac{2\delta}{1-\delta_1^2} \qquad\qquad f\left(\delta \mid \text{return}\right) = \frac{2(1-\delta)}{1-2\delta_1+\delta_1^2}$$

With these building blocks in hand, the firm's revenue from sending a risky product in period 1 with price p_1^r can be derived from adding the revenues from periods 1 and 2, considering the equilibrium of the second-period subgame as describe in Lemma 2.2. Namely:

$$\pi_1^r(p_1^r) = \left(\int_{\delta_1}^{1} s\,ds\right)\Big(p_1^r + \pi_2\big(f\left(\delta \mid \text{keep}\right)\big)\Big) + \left(\int_{\delta_1}^{1} (1-s)\,ds\right)\pi_2\big(f\left(\delta \mid \text{return}\right)\big)$$

The first expression denotes the probability that the customer keeps the product, and the firm receives p_1^r, and subsequently, collects π_2 in period 2 based on the firm's optimal response to observing a "keep" action. The second expression denotes the probability that the customer returns the product and the firm receives 0, and collects π_2 in period 2 based on its optimal response to observing a "return" action.

Since $\pi_1^s = 1$, the firm sends a risky product in period 1 if and only if $\pi_1^r > 1$. This implies that the firm will engage in experimentation with the customer in period 1 if the expected revenue from doing so is higher. Recall that for a given set of prices and realized value (for a risky product), the only determinant of the customer's response is the hassle cost of return, h. The following theorem establishes that an optimal threshold, h^{dp}, exists and is unique such that:

Theorem 2.1. *Under dynamic pricing, $p_1^{s*} = \frac{1}{2}$ and p_1^{r*} is the unique implicit solution to the equation:*

$$\frac{\partial}{\partial p_1^r}\left(\int_{\frac{h}{1-p_1^r+h}}^{1} s p_1^r ds + \int_{\frac{h}{1-p_1^r+h}}^{1} (1-s)\frac{1}{2}ds\right) = 0$$

Furthermore, there exists a unique threshold h^{dp}, such that if $h < h^{dp}$, the firm sends a risky product in period 1 (exploration equilibrium); and if $h \geq h^{dp}$, the firm sends *a safe product in period 1 (exploitation equilibrium).*

All proofs are provided in Appendix A

Theorem 2.1 implies that the firm's optimal revenue only depends on the customer's hassle cost of return.

Under the exploration equilibrium, the firm must compensate the customer for the risk she takes when receiving a risky product by pricing it below her maximum reservation value. We observe that it is optimal for the firm to offer the risky product in period 1 with a lower price such that it induces a customer with a lower δ to join the service. The increase in the period 2 price following a keep compensates the firm for the lower price in period 1 by extracting higher revenue, which we call the *learning effect.* Proposition 2.1 establishes the firm's and the customer's equilibrium strategies under the exploration subgame.

Proposition 2.1. *Under the exploration equilibrium, (a) the firm announces* (p_1^{s*}, p_1^{r*})*; (b) the customer joins the service if her* $\delta \geq \frac{h}{1-p_1^{r*}+h}$*; (c) the firm sends a risky product; and (d) the customer keeps the product if her realized value is* 1 *and returns it, otherwise.*

The second period's subgame equilibrium is as follows:

- *(i) Following a keep observation, the firm's posterior distribution of* δ *on the interval* $[\delta_1, 1]$ *is given by:* $f(\delta) = \frac{2\delta}{1-\delta_1^2}$*. The subgame equilibrium is as follows: (a) the firm announces* (p_2^{s*}, p_2^{r*})*, where* $p_2^{s*} = \frac{1}{2}$*, and* p_2^{r*} *is the unique solution to the implicit equation* $(1 - p_2^r + h)^4 = h^3 (1 + 2p_2^r + h)$*; (b) the customer resumes the service if her* $\delta \geq \frac{h}{1-p_2^{r*}+h}$*; (c) the firm sends a risky product; and (d) the customer keeps the product if her realized value is 1 and returns it, otherwise.*
- *(ii) Following a return observation, the firm's posterior distribution of* δ *on the interval* $[\delta_1, 1]$ *is given by:* $f(\delta) = \frac{2(1-\delta)}{1-2\delta_1+\delta_1^2}$*. The subgame is as follows: (a) the firm announces* (p_2^{s*}, p_2^{r*})*, where* $p_2^{s*} = \frac{1}{2}$*, and* $p_2^{r*} = p_1^{r*}$*; (b) the customer resumes the service; (c) the firm sends a safe product; and (d) the customer keeps the product.*

All proofs are provided in Appendix A

Under the exploitation equilibrium, the firm forgoes the benefits of learning and personalization because it becomes increasingly costly to the firm to lower the price of the risky product to convince the customer to join the service. Proposition 2.2 highlights the firm's and customer's equilibrium strategies under the exploitation subgame.

Proposition 2.2. *Under the exploitation equilibrium, the first period's subgame equilibrium is as follows: (a) the firm announces* (p_1^{s*}, p_1^{r*})*; (b) the customer joins the service; (c) the firm sends a safe product; and (d) the customer keeps the product.*

In the second period, following a keep observation, the firm's posterior belief is that δ *is uniformly distributed on the interval* $[0,1]$*. The second period's subgame equilibrium is as follows: (a) the firm announces* (p_2^{s*}, p_2^{r*})*, where* $p_2^{s*} = \frac{1}{2}$ *and* p_2^{r*} *is the solution to* $h^2(1 + p_2^r + h) = (1 - p_2^r + h)^3$*; (b) the customer resumes the service; (c) the firm sends a safe product; and (d) the customer keeps the product.*

All proofs are provided in Appendix A

3.3 Fixed Pricing Equilibrium

Suppose that the firm commits *ex ante* to a fixed price scheme, which implies $(p_1^s, p_1^r) = (p_2^s, p_2^r)$. We denote the prices with (p^s, p^r) for simplicity[3]. Fixed pricing is particularly relevant in settings where implementing price changes is costly or impractical (Aviv and Pazgal, 2008). Under a fixed pricing mechanism, the firm announces the prices for each product category, (p^r, p^s), at the start of period 1. The customer takes this announcement as given (i.e. fixed) for both subscription periods, and decides whether to join the service. If she joins, the firm must then choose from which category to send a product. The customer observes her value for the product and opts to keep or return it. Having observed the customer's first period decisions, the firm

---

[3] Here we consider the case where the firm announces prices such that $\frac{1}{2} < p^r \leq 1$ and $p^s \leq \frac{1}{2}$, so that either categories are viable options to be sent to the customer and the customer joins/resumes the service. Intuitively, if $p^r \leq \frac{1}{2}$ it would never be optimal for the firm to send a risky product since the expected revenue would be less than offering the safe category priced at $\frac{1}{2}$. Meanwhile if the firm prices the safe (risky) category at $p^s > \frac{1}{2}$ ($p^r > 1$) the customer does not participate in the game and the firm is left with 0 payoff. The validity of this statement is proven in the proof of Lemma 2.3.

then updates its belief about the customer's taste, and decides from which category to send a product in the second period. As with the dynamic pricing mechanism, we analyze the two-period stochastic game between the firm and the customer in reverse chronological order, while seeking pure-strategy subgame perfect Bayesian equilibria.

Second-period subgame: Assume in period 1 the customer with a $\delta \geq \delta_1$ has joined the service. In period 2's subgame, the customer's decision to keep (return) a product, and to resume (leave) the service, is identical to the dynamic pricing case. Following the rational expectations framework, if the customer believes that upon resuming the service the firm will send a risky product (which is correct in equilibrium), she chooses to resume if her $\delta \geq \frac{h}{1-p^r+h}$. On the other hand, if she believes that upon resuming the service the firm will send a safe product, she chooses to resume if $p^s \leq \frac{1}{2}$. Since the prices are fixed, the firm's decision in period 2 is limited to the product choice. The firm sends a risky product if the expected revenue from sending a risky product is greater than the expected revenue from sending a safe product. the expected revenues from sending a safe and risky product are given by:

$$\pi_2^s = p^s \tag{2.3a}$$

$$\pi_2^r = p^r \int\limits_{\delta_2}^{1} sf\left(s\right)ds \tag{2.3b}$$

where $f\left(\cdot\right)$ is the posterior pdf of the firm's updated belief about δ, after observing the customer's action in the first period and $\delta_2 = \max\left(\frac{h}{1-p^r+h}, \delta_1\right)$.

Note that unlike the case of dynamic pricing, since period 2's prices are not variables on which the firm can optimize in the second period, it implies that in solving for the optimal revenue in period 2, δ_2 is treated as a fixed parameter. Therefore, to simplify our notation, we may denote the expected probability of the customer keeping the product in the second period given $f(\cdot)$ as $\delta^u = \int_{\delta_2}^{1} sf\left(s\right)ds$. This brings us to the following lemma.

Lemma 2.3. *Given any δ_1, posterior pdf $f(\cdot)$, p^s and p^r, there exists a unique subgame perfect equilibrium in the second period played by the firm and the customer as follows:*

(i) If $\delta^u p^r > p^s$, then (a) the customer resumes the service if her $\delta \geq \frac{h}{1-p^r+h}$; (b) the firm sends a risky product to the customer; (c) the customer keeps the product if her realized value is $v = 1$, and returns it, otherwise.

(ii) If $\delta^u p^r \leq p^s$, then (a) the customer resumes the service; (b) the firm sends a safe product; (c) the customer keeps the product.

All proofs are provided in Appendix B.

Considering period 2's equilibrium as described in Lemma 2.3, if the firm's optimal strategy is to send a risky product, the customer with $\delta > \delta_2$ is left with a surplus of $\delta(1-p^r) - (1-\delta)h$. Meanwhile, if it is optimal for the firm to send a safe product, the customer is left with a surplus $\frac{1}{2} - p^s$.

First-period subgame: The analysis of this section considers the strategic joining behavior of the customer, and is centered around two main questions. The first pertains to the optimal prices: how should the firm optimally make its pricing decisions? The second question concerns the firm's equilibrium product choice: given the optimal prices of the safe and risky products, is it beneficial for the firm to experiment with the customer? Theorem 2.2 describes the unique threshold-type pure strategy equilibrium of the game, such that for h values below a threshold the firm finds learning to be beneficial.

Theorem 2.2. *Under fixed pricing, there exists a unique threshold h^{fp}, such that if $h < h^{fp}$, the firm sends a risky product in period 1 (exploration equilibrium). Meanwhile, if $h \geq h^{fp}$, the firm sends a safe product in period 1 (exploitation equilibrium).*

All proofs are provided in Appendix B

Under the exploration equilibrium, the firm sets the price of the safe product to $p^{s*} = \frac{1}{2} - h$. With this p^{s*}, the customer joins the service regardless of her δ (i.e.

$\delta_1 = 0$), and receives a risky product in period 1. If she likes the risky product, she keeps it. That said, she may choose to either resume or leave the service, depending on whether her expected utility from receiving another risky product in period 2 is non-negative. In other words, since $\delta_1 = 0$, a customer with an unfavorable taste for risky products (i.e. $\delta \ll 1$) still chooses to join, and with a small probability, may like the risky product and keep it. Given that the firm's optimal choice upon observing a "keep" action is to send a risky product in the second period, the customer expecting a negative utility in the second period (i.e. $U_2(\delta) = \delta(1 - p^r) - (1 - \delta)h < 0$), leaves. However, if the customer dislikes the risky product, she returns it and receives a safe product in period 2 at a price below her reservation value ($v = \frac{1}{2}$). Note that with this pricing policy, the firm discounts the price of the safe product category for the benefit of making a more informed product choice decision in period 2, without compromising the price of the risky product category. The $h < h^{fp}$ threshold from Theorem 2.2 implies that this is only beneficial if the firm can maintain its guaranteed revenue from the safe category. In other words, if the firm needs to lower the price of safe products due to the high hassle cost of return, then it becomes overly costly for the firm to send a risky product (regardless of its expectation on the customer) because the customer may return the product, and the firm will have to send a low-priced safe product in response. The result also implies that the firm benefits from the presence of the strategic customer. Namely, the customer with an unfavorable δ still joins the service expecting that she will be compensated for her participation in the experimentation phase in period 1, by enjoying a positive surplus in period 2.

Proposition 2.3. *Exploration equilibrium: If* $h < h^{fp}$*, there exists a unique period 1 subgame equilibrium such that (a) the firm announces* (p^{s*}, p^{r*})*, where* $p^{s*} = \frac{1}{2} - h$ *and* p^{r*} *is the implicit solution to* $5\left(1 - p^r + h\right)^4 = 2h^3\left(1 + 2p^r + h\right)$*; (b) the customer joins the service regardless of her* δ*; (c) the firm sends a risky product in period 1; and (d) the customer keeps it if her realized value is* $v = 1$*, and returns it, otherwise.*

Furthermore,

(i) Following a keep observation, the firm's posterior distribution of δ *has a* $f(\delta) = 2\delta$*. (a) The customer resumes the service if her* $\delta \geq \frac{h}{1-p^{r*}+h}$*, (b) the firm sends a risky product, and (c) the customer keeps the product if her realized value is* $v = 1$*, and returns it, otherwise.*

(ii) Following a return observation, the firm's posterior distribution of δ *has a* $f(\delta) = 2(1-\delta)$*. (a) The customer resumes the service, (b) the firm sends a safe product, and (c) the customer keeps the product.*

Similar to the dynamic pricing mechanism, in a market with $h \geq h^{fp}$, the cost of learning outweighs its benefit and hence, the firm opts to send a safe product to the customer and avoids experimentation.

Proposition 2.4. *Exploitation equilibrium: If* $h \geq h^{fp}$*, there exists a unique period 1 subgame equilibrium such that: (a) the firm sets* (p^{s*}, p^{r*})*, where* $p^{s*} = \frac{1}{2}$ *and* p^{r*} *is the implicit solution to* $(1 - p^r + h)^3 = h^2(1 + p^r + h)$*; (b) the customer joins the service; (c) the firm sends a safe product in period 1; (d) the customer keeps the product.*

Furthermore, in the second period, the customer resumes the service and the firm sends a safe product in period 2, which the customer keeps.

4 Discussion

We discuss several important implications of the results we obtained from the previous section. First, we compare the dynamic and fixed pricing equilibria in terms of the firm's pricing strategies and its expected revenue. We further investigate the implications of each pricing mechanism on the firm's customer acquisition and retention capability. Second, we focus on the implications of different return strategies for the firm. Finally, we discuss the implication of our results for consumer surplus and social welfare.

4.1 Dynamic vs. Fixed Pricing

We first examine whether dynamic pricing is preferred by the firm, in the presence of personalization. Intuitively, dynamic pricing should be preferred by the firm since the flexibility offered by dynamic pricing enables the firm to take advantage of the learning that takes place in period 1. In other words, whereas with a fixed pricing strategy the firm's only lever in period 2's subgame is the product choice, it is expected that the firm would benefit from a dynamic pricing strategy since it offers an additional degree of freedom, temporal pricing. Interestingly however, we observe that the opposite is true. Proposition 2.5 describes this observation, which is illustrated in Fig. 2·1.

Proposition 2.5. *In the presence of experimentation, the firm's expected revenue is higher under fixed pricing than it is under dynamic pricing.*

All proofs are provided in Appendix C.

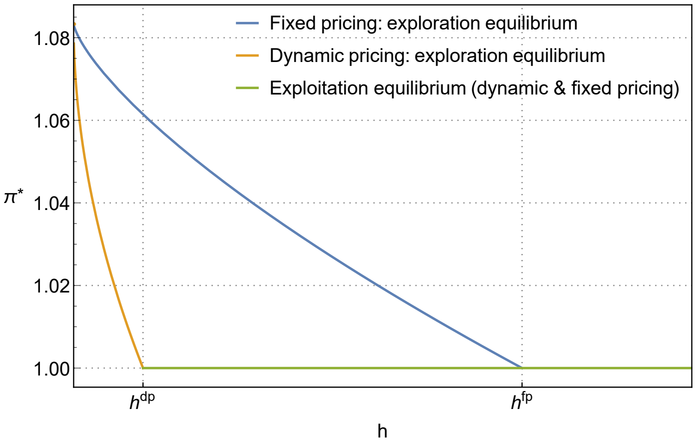

Figure 2·1: Optimal revenue as a function of h under dynamic and fixed pricing.

According to Fig. 2·1, fixed pricing yields a higher revenue than dynamic pricing. Furthermore, the firm can experiment with greater values of hassle cost ($h^{fp} > h^{dp}$). The underlying mechanism that favors fixed pricing over dynamic pricing arises from the customer's forward-looking strategic behavior. Note that the strategic customer joins the service if her expected utility is non-negative. Under fixed pricing, the firm discounts the price of the safe product so that a customer with an unfavorable taste for the risky product participates in period 1's experiment, knowing that she will receive a safe product with a discounted price in period 2. Under dynamic pricing, for $0 < h < h^{dp}$, the threshold above which the customer joins, δ_1, is strictly greater than 0 and therefore, the firm loses the customer at the initial point of interaction if her $\delta < \delta_1$. This occurs with a probability of δ_1. However, under fixed pricing, the customer joins regardless of her δ since $\delta_1 = 0$, which ensures full acquisition at the point of contact.

Furthermore, as illustrated in Fig. 2·2, the firm's single optimal price for products from the risky category under fixed pricing, is greater than those of each period under dynamic pricing (i.e. $p_1^{dp}, p_2^{dp} < p^{fp}$). The rationale behind this counterintuitive observation is related to the discounted price of the safe category, which is depicted in Fig. 2·3. Under fixed pricing, by pricing the safe category below the customer's reservation price, the firm is able to increase the price of products from the risky category without compromising customer acquisition. Namely, the firm does not face the risk that a subscription might not occur at all, in which case, the game would end with the firm earning zero revenue.

To better highlight the differences between the firm's exploration equilibrium under fixed and dynamic pricing, we turn to the customer's behavior with respect to joining (and staying) in the subscription service. Fig. 2·5 presents the firm's acquisition threshold, δ_1, and retention threshold, δ_2, under fixed and dynamic pricing as a function of hassle cost. As illustrated in the figure, the higher price of products from the risky category under fixed pricing results in a higher δ_2^{fp}, and subsequently, a lower retention in period 2. However, since $\delta_1^{fp} = 0$, the firm not only enjoys full acquisition in period 1, but also the additional expected revenue generated by a low-δ customer who may like the risky product in period 1 with probability δ, and choose to keep it, but leave the subscription in period 2.

4.2 Optimal Return Policy

The next question we investigate is whether the firm should allow returns. Let us consider the period 2 subgame when the firm enforces a no-return policy. When returns are not allowed, the customer's expected utility from receiving a risky product is $U^r(\delta) = \delta \times 1 - p_2^r$; therefore she resumes the service if $p_2^r \leq \delta$. The firm's revenue from sending a risky product equals the probability that the customer joins the service (i.e. $1 - F(p_2^r) = 1 - p_2^r$) and pays the price p_2^r: $\pi(p_2^r) = (1 - p_2^r)\, p_2^r$, which for any p_2^r, is less than the revenue generated from sending a safe product, $\frac{1}{2}$. This case is repeated in period 1's subgame. Consequently, when the firm enforces a no-return policy, it is never optimal to send a risky product to the customer.

Therefore, a firm in a market with low hassle costs (i.e. where experimentation is optimal) is better off if it allows returns with full refunds, which allows it to learn customer tastes through experimentation. Meanwhile, in a market with high hassle cost of returns (i.e. where exploitation is optimal), a firm is indifferent between not allowing/allowing returns as it only offers safe products, for which the customer has a known (and fixed) value, and always keeps.

Proposition 2.6. *For any values of the hassle cost, the firm's expected revenue is weakly greater under a full-refund return policy compared to a no-return policy.*

All proofs are provided in Appendix C.

Under the exploration equilibrium, we find that there are two main mechanisms through which the firm benefits from allowing returns. The first mechanism echos a recurring theme in the recent literature that considers the benefits of refund policies as an uncertainty resolution mechanism (Cachon and Feldman, 2018). In the presence of strategic customers, a full-refund return policy allows the firm to attract customers at higher prices. The second mechanism is by learning the customer's taste, which is unique to the subscription business model of personalization, and occurs through repeated interactions with the customer.

An important byproduct of Theorems 2.1 and 2.2 is that under both fixed and dynamic pricing, the firm's revenue is maximized when the exogenous hassle cost of return is zero (i.e. $h = 0$). In this case, the customer does not incur any return cost when faced with a product they don't like. An important question that remains is whether or not the firm should share the burden of returns if $h > 0$. In practice, there are several major curated subscription box businesses, most notably Nordstrom's TruckClub, that provide hassle-free returns through a 24/7 customer service and free at-home pickups.

We modify the model presented in section 3 such that given a non-zero exogenous h, the firm can choose to incur a portion of the cost. We assume h is split between the customer and the provider: the customer incurs $(1-\alpha)h$, and the provider incurs αh. Therefore, if the customer decides to keep a product of value v at price p, her utility is $u = v - p$. If she decides to return the product, her utility is $-(1-\alpha)h$. We find that under both fixed and dynamic pricing policies, the firm is strictly better off if it takes all the hassle cost to itself (i.e. $\alpha = 1$). Not only can the firm generate more revenue by providing hassle-free returns, but also, it can experiment with greater values of h. In other words, $h^{\alpha} > h^{fp} > h^{dp}$, where h^{α} is the hassle cost threshold below which experimentation is optimal. This is elaborated in Lemma 2.4.

Lemma 2.4. *With an exogenous $h < h^{\alpha}$, the firm's unique optimal return policy is such that $\alpha^* = 1$, $p_1^{r^*} = p_2^{r^*} = 1$, and $p_1^{s^*} = p_2^{s^*} = \frac{1}{2}$ under both fixed and dynamic pricing policies. The firm's unique optimal product choice is such that:*

(i) In period 1, the firm sends a risky product.

(ii) In period 2, the firm sends a risky product if it observes "keep" in period 1 and sends a safe product, otherwise.

All proofs are in Appendix C.

Fig. 2·6 illustrates the numerical comparison of the firm's revenue when $\alpha^* = 1$, and when $\alpha = 0$ (i.e. the base model). There are several mechanisms through which the profit function is affected by the introduction of α. First, taking away the hassle cost of returns allows the firm to charge a higher price for the risky product. Second, with $\alpha = 1$ the customer faces no disutility from joining/resuming the service, allowing the firm to attract and keep more customers. Noteworthy is that the *disadvantage* of $\alpha = 1$ is not only that the firm incurs h, but it is also detrimental to the firm's learning. This is because the customer joins/resumes the service regardless of her δ when $(1 - \alpha)\, h = 0$. Therefore, there is no truncation of the state space (i.e. customers) from which the firm can learn. Nonetheless, the overall outcome of internalizing all of the hassle cost is beneficial to the firm.

4.3 Implications for Social Welfare

Social welfare (SW) is the sum of customer surplus (CS) and the firm's revenue $SW = CS + \pi$. Therefore, to study the implications of our results for social welfare, we must consider the customer' perspective. We do so for three pricing strategies: (i) dynamic pricing with $\alpha = 0$, (ii) fixed pricing with $\alpha = 0$, and (iii) optimal cost sharing (where $\alpha^* = 1$ and there is no difference between dynamic and fixed pricing strategies).

In all three pricing strategies, note that $CS = 0$ and $SW = 1$ in the exploitation equilibrium since $v = p = \frac{1}{2}$. Therefore, differences in the exploration equilibrium will determine the differences in CS and SW. Under dynamic pricing with $\alpha = 0$, for $h < h^{dp}$, the customer expects a positive surplus from receiving a risky product if her $\delta > \delta_1$. If $\delta_1 \leq \delta < \delta_2$, given the equilibrium path described in lemma 2.1, the customer's expected utility is $U_1(\delta) = \delta\,(1 - p_1^r) + (1 - \delta)\left(-h + \frac{1}{2} - p_2^s\right)$. Namely, the customer receives a risky product in period 1; if she likes the product with probability δ, she keeps it and leaves the service; and if she dislikes the product with probability $1-\delta$, she returns it and resumes the service to receive a safe product in period 2. Meanwhile, if $\delta_2 \leq \delta \leq 1$, the customer's expected utility is $U_1(\delta) = \delta\left(1 - p_1^r + \delta(1-p_2^r) - (1-\delta)h\right) + (1-\delta)\left(-h + \frac{1}{2} - p_2^s\right)$. The customer receives a risky product in period 1. She keeps the product and resumes the service to receive a risky product in period 2 with probability δ; and returns the product and resumes the service to receive a safe product in period 2 with probability $1-\delta$. Therefore, under dynamic pricing, the total expected customer surplus is:

$$\begin{aligned} CS^{dp} = &\int_{\delta_1}^{\delta_2} \left(s\left(1-p_1^{r*}\right) - (1-s)h\right) ds + \\ &\int_{\delta_2}^{1} \left(s\left(1-p_1^{r*}\right) - (1-s)h + s\left(s\left(1-p_2^{r*}\right) - (1-s)h\right)\right) ds \end{aligned}$$

Under fixed pricing with $\alpha = 0$, for $h < h^{fp}$, the customer expects a positive surplus regardless of her δ, such that if her $0 \leq \delta < \delta_2$, given the equilibrium path described in lemma 2.3, $U_1(\delta) = \delta\left(1-p^r\right) + (1-\delta)\left(-h + \frac{1}{2} - p^s\right)$ and if her $\delta_2 \leq \delta \leq 1$, $U_1(\delta) = \delta\left(1 - p^r + \delta(1-p^r) - (1-\delta)h\right) + (1-\delta)\left(-h + \frac{1}{2} - p^s\right)$. Replacing the parameters with their equilibrium values, the total expected customer surplus is as follows:

$$CS^{fp} = \int_{0}^{\delta_2} s\left(1-p^{r*}\right) ds + \int_{\delta_2}^{1} \left(s\left(1-p^{r*}\right) + s\left(s\left(1-p^{r*}\right) - (1-s)h\right)\right) ds \tag{2.4}$$

Finally, under hassle cost sharing ($\alpha^* = 1$), for $h < h^{\alpha}$, $\pi = \frac{13}{12} - \frac{2}{3}h$, $CS = 0$ and $SW = \frac{13}{12} - \frac{2}{3}h$. Note that when $\alpha = 1$, the firm prices the safe category at $\frac{1}{2}$, and the risky category at 1, which extracts all surplus.

Fig. 2·7 and Fig. 2·8 provide equilibrium values of the total consumer surplus and social welfare. As illustrated in this figure, under the exploration equilibrium, if the customer is the party who incurs the hassle cost of return, in aggregate, the customer is left with some surplus, which is increasing in h. The reason is that the firm's optimal prices under both dynamic and fixed pricing are strictly *decreasing* in h. However, if the firm is the party that incurs the hassle cost of return, the firm extracts all surplus through perfect price discrimination. Therefore, surprisingly, customers are better off if their hassle cost of return is non-zero. Additionally, it can be observed that social welfare is higher under the exploration subgame equilibrium of both dynamic and fixed pricing, which further underscores the expected value generated by personalization.

5 Conclusion

We study how a subscription service firm should utilize its operational levers to learn a strategic customer's taste and personalize its services to her preferences. We find that depending on the hassle cost of receiving an unwanted product, it can be optimal for a firm to experiment with its content. Furthermore, we show that when facing forward-looking customers, a fixed pricing strategy not only yields higher expected revenue and allows the firm to experiment with higher hassle costs of return, but also, maximizes customer acquisition. Fixed pricing also benefits the customer since it leaves her with non-zero surplus.

Clearly, personalization can only be effective if the cost of experimentation (sending unwanted products to the customer) is not excessive. Furthermore, fixed pricing is the optimal pricing strategy only if the customers are forward-looking. Namely, the customer with an unfavorable taste for risky products is willing to participate in the experimentation with the knowledge that if they are patient, they will receive their preferred product at a discounted price once the firm learns their preferences.

Interestingly, we find that a hassle-free return is desirable for the firm even if it requires the firm to internalize the hassle cost to make this possible for the customer. This is because the firm can offer products at the customer's reservation price, while experimenting with risky, yet of higher value, products. The hassle-free returns strategy highlights the benefits of learning in a setting with repeated interactions with the customer, even if it obligates the firm to not only fully refund the customer but also, make it very easy for her to return unwanted products.

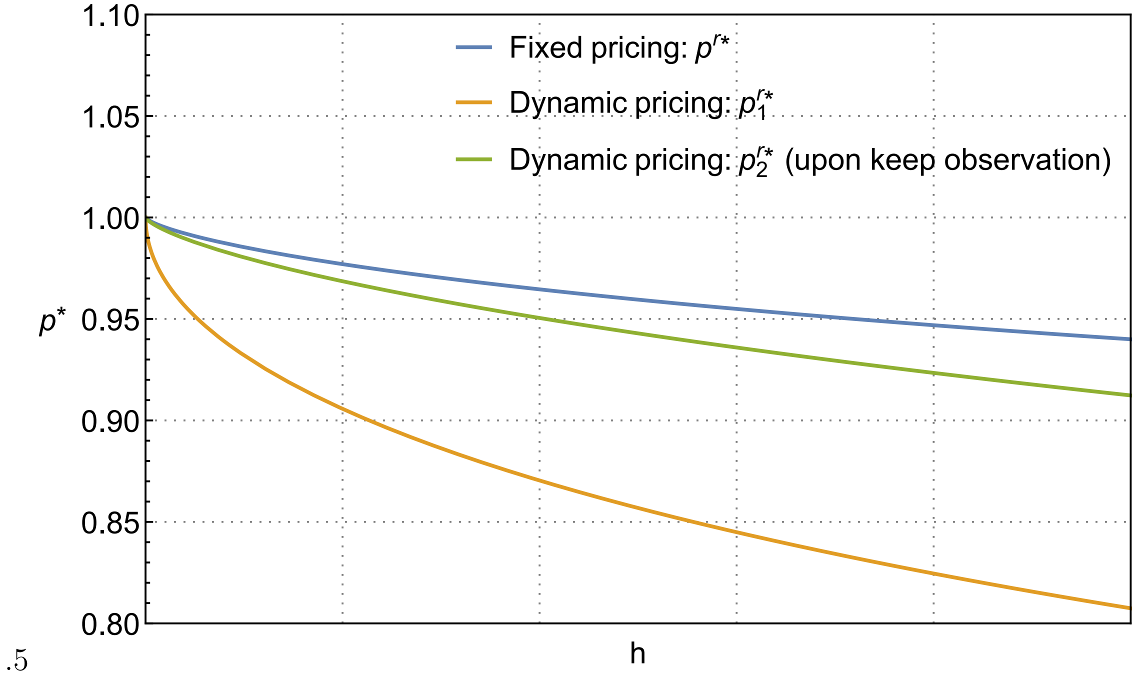

Figure 2·2: Risky category

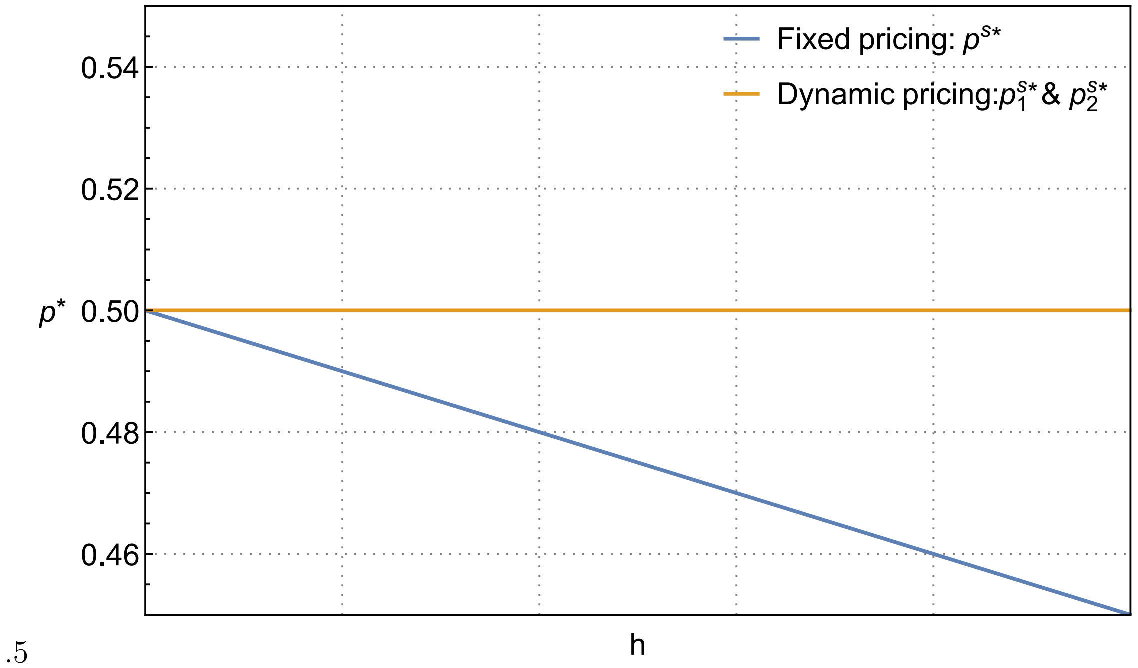

Figure 2·3: Safe category

Figure 2·4: Optimal Prices under exploration equilibrium of dynamic and fixed pricing

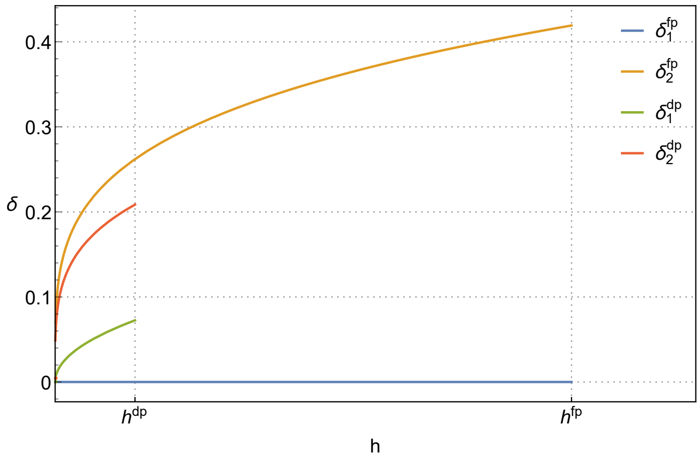

Figure 2·5: Acquisition and retention under exploration equilibrium of dynamic and fixed pricing.

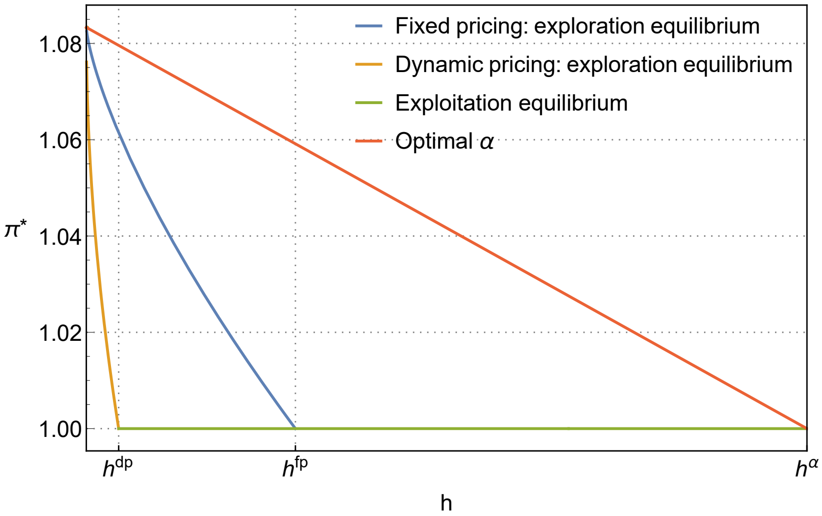

Figure 2·6: Optimal revenue as a function of h under dynamic and fixed pricing with $\alpha = 0$, and $\alpha^* = 1$

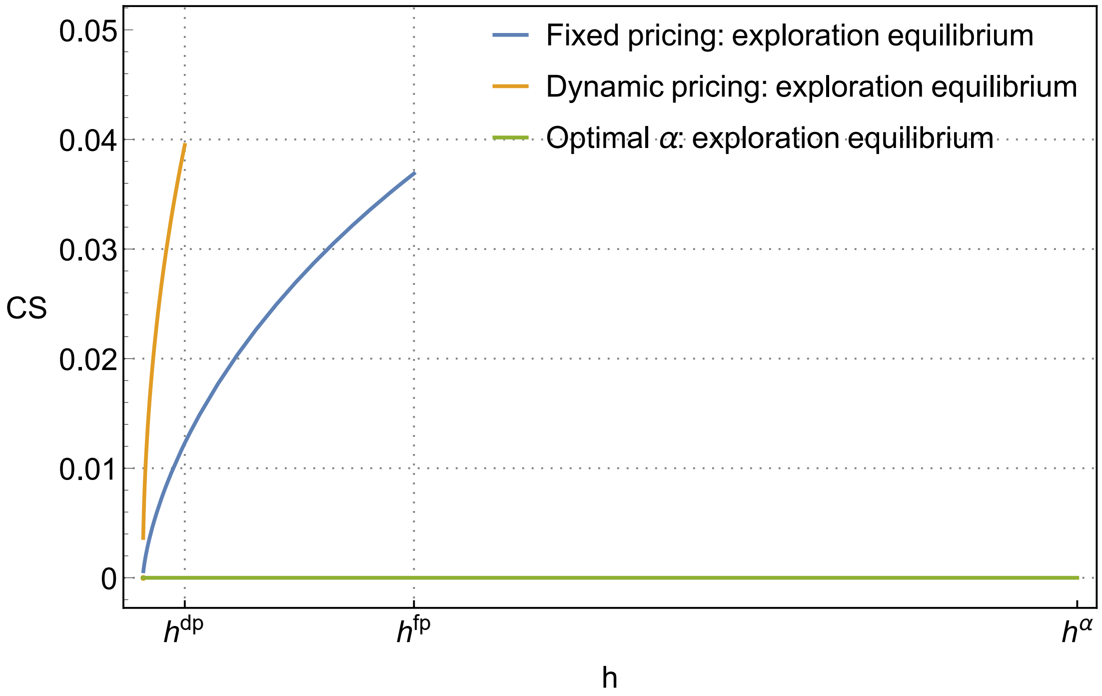

.5

Figure 2·7: Consumer surplus

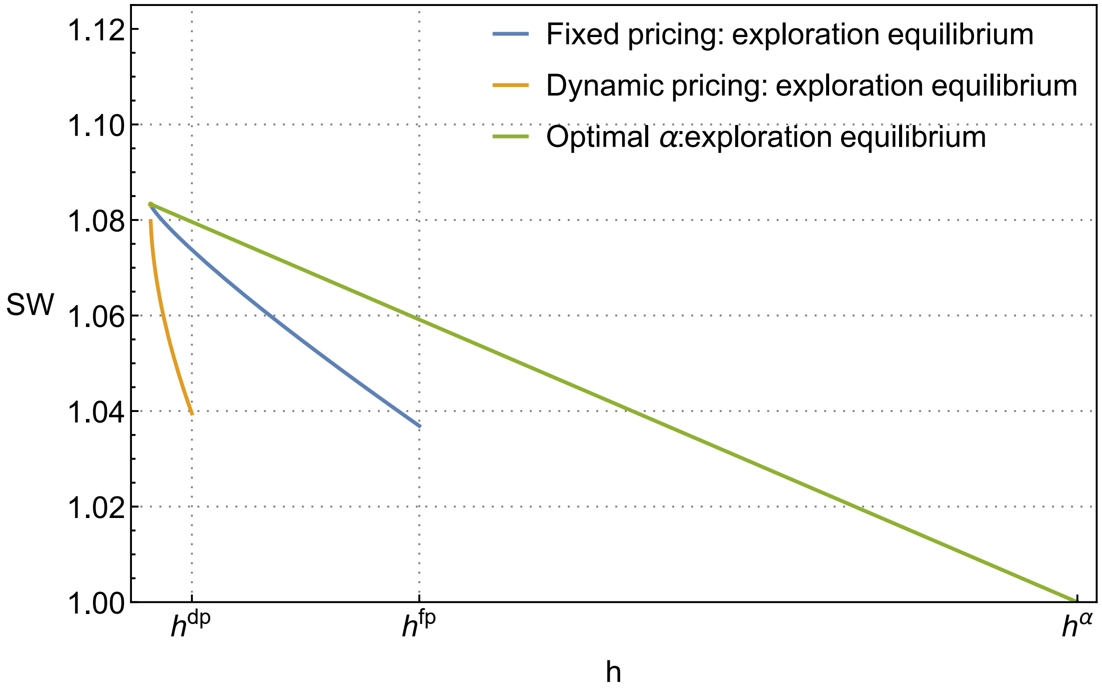

.5

Figure 2·8: Social welfare

Figure 2·9: Expected consumer surplus and social welfare as a function of h under dynamic and fixed pricing with $\alpha = 0$ and optimal α

References

Altug, M. S. and Aydinliyim, T. (2016). Counteracting strategic purchase deferrals: The impact of online retailers' return policy decisions. *Manufacturing & Service Operations Management*, 18(3):376–392.

Aviv, Y. and Pazgal, A. (2008). Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing & service operations management*, 10(3):339–359.

Aviv, Y. and Vulcano, G. (2012). Dynamic list pricing. In *The Oxford handbook of pricing management*.

Aydin, G. and Ziya, S. (2009). Personalized dynamic pricing of limited inventories. *Operations Research*, 57(6):1523–1531.

Ban, G.-Y. and Keskin, N. B. (2020). Personalized dynamic pricing with machine learning: High dimensional features and heterogeneous elasticity. *Available at SSRN 2972985*.

Bernstein, F., Chakraborty, S., and Swinney, R. (2020). Dynamic content variation with customer learning. *Available at SSRN 3489425*.

Bernstein, F., Kök, A. G., and Xie, L. (2015). Dynamic assortment customization with limited inventories. *Manufacturing & Service Operations Management*, 17(4):538–553.

Bernstein, F., Modaresi, S., and Sauré, D. (2019). A dynamic clustering approach to data-driven assortment personalization. *Management Science*, 65(5):2095–2115.

Cachon, G. P. and Feldman, P. (2018). Pricing capacity over time and recourse strategies: Facilitate reselling, offer refunds/options, or overbook? *Available at http://dx.doi.org/10.2139/ssrn/3235757.*

Cachon, G. P. and Swinney, R. (2009). Purchasing, pricing, and quick response in the presence of strategic consumers. *Management Science*, 55(3):497–511.

Candogan, O., Bimpikis, K., and Ozdaglar, A. (2012). Optimal pricing in networks with externalities. *Operations Research*, 60(4):883–905.

Caro, F. and Gallien, J. (2007). Dynamic assortment with demand learning for seasonal consumer goods. *Management Science*, 53(2):276–292.

Chen, T., Fenyo, K., Yang, S., and Zhang, J. (2018). Thinking inside the subscription box: New research on e-commerce consumers. *McKinsey & Company*, pages 1–9.

Cheung, W. C. and Simchi-Levi, D. (2017). Thompson sampling for online personalized assortment optimization problems with multinomial logit choice models. *Available at SSRN 3075658.*

DeGroot, M. H. (2005). *Optimal statistical decisions*, volume 82. John Wiley & Sons.

Feldman, P., Papanastasiou, Y., and Segev, E. (2019). Social learning and the design of new experience goods. *Management Science*, 65(4):1502–1519.

Hu, X., Wan, Z., and Murthy, N. N. (2019). Dynamic pricing of limited inventories with product returns. *Manufacturing & Service Operations Management*, 21(3):501–518.

Li, K. J. (2021). Behavior-based quality discrimination. *Manufacturing & Service Operations Management*, 23(2):425–436.

Mišić, V. V. and Perakis, G. (2020). Data analytics in operations management: A review. *Manufacturing & Service Operations Management*, 22(1):158–169.

Murthi, B. and Sarkar, S. (2003). The role of the management sciences in research on personalization. *Management Science*, 49(10):1344–1362.

Papanastasiou, Y. and Savva, N. (2017). Dynamic pricing in the presence of social learning and strategic consumers. *Management Science*, 63(4):919–939.

Rusmevichientong, P., Shen, Z.-J. M., and Shmoys, D. B. (2010). Dynamic assortment optimization with a multinomial logit choice model and capacity constraint. *Operations research*, 58(6):1666–1680.

Shen, Q. and Miguel Villas-Boas, J. (2018). Behavior-based advertising. *Management Science*, 64(5):2047–2064.

Strauss, A. K., Klein, R., and Steinhardt, C. (2018). A review of choice-based revenue management: Theory and methods. *European Journal of Operational Research*, 271(2):375–387.

Su, X. (2010). Intertemporal pricing and consumer stockpiling. *Operations research*, 58(4-part-2):1133–1147.

Su, X. and Zhang, F. (2008). Strategic customer behavior, commitment, and supply chain performance. *Management Science*, 54(10):1759–1773.

Appendix

Appendix A: Dynamic Pricing- General Case

Proof of Lemma 2.2.

Proof. Suppose the firm believes for an arbitrary δ_1 that the customer joins the service if her $\delta \geq \delta_1$. If the firm sends a safe product in period 1, as long as p_1^s is such that $\frac{1}{2} - p_1^s \geq -h$, the customer keeps the product. Thus, when $p_1^s \leq \frac{1}{2} + h$ the firm's posterior belief of δ remains unchanged (i.e. $f(\delta) \sim U[\delta_1, 1]$).

In the second period, expecting a safe product, the customer resumes the service if $p_2^s \leq \frac{1}{2}$, and keeps the product. Given $f(\delta)$, the firm's expected revenue from sending a safe product equals $\pi_2^s = p_2^s$. The expected revenue is increasing in p_2^s. Therefore, the firm's optimal price for the safe product in the second period is unique and is $p_2^{s*} = \frac{1}{2}$. In the second period, if the firm opts to send a safe product priced at $p_2^{s*} = \frac{1}{2}$, the customer resumes the service, receives the product and keeps it. Consequently, $\pi_2^{s*} = \frac{1}{2}$.

Let us consider the subgame where given the updated belief $f(\delta)$, the firm opts to send a risky product. If $p_2^r \leq h$, the customer resumes the service if her $\delta \geq p_2^r$ and keeps the product regardless of her realized value. Therefore, the firm's expected revenue equals $\pi_2^r = \int_{\delta_2}^1 p_2^r$, where $\delta_2 = \max(\delta_1, p_2^r)$. Given our assumption that $0 \leq h \leq \frac{1}{2}$, when $p_2^r \leq h$, $\pi_2^r \leq \frac{1}{2}$ for any δ_2. Consequently, it is never optimal for the firm to send a risky product priced at $p_2^r \leq h$.

If $p_2^r > h$, the customer keeps a risky product if her realized value is 1 and returns it otherwise. Moreover, the customer resumes the service if her δ is such that $\delta(1 - p_2^r) - (1 - \delta)h \geq 0 \rightarrow \delta \geq \frac{h}{1 - p_2^r + h}$. Therefore, the firm's expected revenue from sending a risky product equals $\pi_2^r = \int_{\delta_2}^1 p_2^r s f(s)\, ds$, where, $\delta_2 = \max\left(\frac{h}{1-p_2^r+h}, \delta_1\right)$. If $\frac{h}{1-p_2^r+h} > \delta_1 \rightarrow p_2^r > 1 - \frac{1-\delta_1}{\delta_1}h$, $\pi_2^r = p_2^r\left(\frac{1}{2} - \frac{h^2}{2\left(1-p_2^r+h\right)^2}\right)$. The first order gives $\frac{1}{2} - \frac{h^2}{2(1-p_2^r+h)^2} - \frac{h^2 p_2^r}{(1-p_2^r+h)^3} = 0$. While the second order is $-\frac{2h^2}{(1-p_2^r+h)^3} - \frac{3h^2 p_2^r}{(1-p_2^r+h)^4} < 0$. Therefore, the expected revenue is concave in p_2^r. If δ_1 is such that $2(1+h)\delta_1^3 - \delta_1^2 h - h > 0$, the revenue function is decreasing in p_2^r and the optimal p_2^{r*} is the lower bound on p_2^r which is $1 - \frac{1-\delta_1}{\delta_1}h$. Otherwise, the revenue function has a unique local maxima at p_2^{r*}, where p_2^{r*} satisfies $h^2(1 + p_2^r + h) = (1 - p_2^r + h)^3$. Meanwhile, if $\frac{h}{1-p_2^r+h} \leq \delta_1 \rightarrow p_2^r \leq 1 - \frac{1-\delta_1}{\delta_1}h$, $\pi_2^r = \frac{1}{2}p_2^r(1 - \delta_1^2)$, which is strictly increasing in p_2^r. Therefore, p_2^{r*} binds the upper constraint on p_2^r: $p_2^r \leq 1 - \frac{1-\delta_1}{\delta_1}h$: $p_2^{r*} = 1 - \frac{1-\delta_1}{\delta_1}h$. As such for a δ_1, there exists a unique $p_2^{r*}(\delta_1)$. The firm sends a risky product priced at $p_2^{r*}(\delta_1)$, in period 2, if and only if the expected revenue from sending a risky product is greater than that of sending a safe product.

We now focus on the second period subgame which follows the firm sending a risky product in the first period. Note, if $h < p_1^r \leq 1 + h$, the customer keeps a risky product if her realized value is 1 and returns it if her realized value is 0. If in the first period the customer keeps a risky product, the firm's updated belief has the following posterior pdf on the interval $[\delta_1, 1]$:

$$f(\delta \mid \text{keep}) = \frac{p(v = 1 \mid \delta) f(\delta)\, d\delta}{\int_{\delta_1}^{1} \delta f(\delta)\, d\delta} = \frac{2\delta}{1 - \delta_1^2}$$

Given $f(\delta)$, the firm's expected revenue from sending a risky product in the second period, is given by:

$$\pi_2^r = \int_{\delta_2}^{1} s p_2^r f(s)\, ds = p_2^r \left(\int_{\delta_2}^{1} \frac{2s^2}{1 - \delta_1^2} ds \right) \tag{2.5}$$

where $\delta_2 = \max\left(\frac{h}{1 - p_2^r + h}, \delta_1\right)$. If $\frac{h}{1 - p_2^r + h} > \delta_1 \rightarrow p_2^r > 1 - \frac{1 - \delta_1}{\delta_1} h$, the firm maximizes:

$$\max_{p_2^r} \left(p_2^r \frac{2\left(1 - \frac{h^3}{\left(1 - p_2^r + h\right)^3}\right)}{3\left(1 - \delta_1^2\right)} \right) \tag{2.6}$$

$$s.t. \quad \max\left(1 - \frac{1 - \delta_1}{\delta_1} h, h\right) < p_2^r \leq 1 \tag{2.7}$$

Checking for the SOC, we have:

$$\frac{d^2 \pi_2^r}{d p_2^{r^2}} = \frac{2}{3\left(1 - \delta_1^2\right)} \left(-\frac{6h^3}{(1 - p_2^r + h)^4} - \frac{12 h^3 p_2^r}{(1 - p_2^r + h)^5} \right) < 0$$

Therefore, the expected revenue is concave in p_2^r. The FOC gives:

$$\frac{d\pi_2^r}{d p_2^r} = \frac{2}{3\left(1 - \delta_1^2\right)} \left(1 - \frac{h^3}{(1 - p_2^r + h)^3} - \frac{3 h^3 p_2^r}{(1 - p_2^r + h)^4} \right)$$

Simplifying and rearranging the terms gives:

$$\frac{d\pi_2^r}{dp_2^r} = \frac{2}{3\left(1-\delta_1^2\right)}\left(1 - \frac{h^3(1+2p_2^r+h)}{(1-p_2^r+h)^4}\right)$$

Let us denote $1 - p_2^r + h = A$, then the expression becomes:

$$\frac{2}{3\left(1-\delta_1^2\right)}\left(\frac{A^4 - h^3\left(3+3h-2A\right)}{A^4}\right) = \frac{2}{3\left(1-\delta_1^2\right)A^4}\left(A^4 + 2Ah^3 - 3h^3\left(1+h\right)\right)$$

Note, given SOC, we know FOC is decreasing in p_2^r (increasing in A). Therefore, if at the maximum value that A can get, the expression is negative, then the expected revenue is decreasing in p_2^r. However, if at the maximum value that A can get, the expression is positive and at the minimum the expression is negative, there exists an interior A that satisfies the first-order condition. Given that $\left(1 - \frac{1-\delta_1}{\delta_1}h, h\right) < p_2^r \leq 1$, the maximum value that A can get equals $A = \frac{h}{\delta_1}$. Replacing A's maximum value and rearranging, the first-order at A^{max} becomes:

$$\frac{2}{3\left(1-\delta_1^2\right)}\left(1 + 2\delta_1^3 - \frac{3\delta_1^4(1+h)}{h}\right)$$

which is negative if $3(1+h)\delta_1^4 > h(1+2\delta_1^3)$ and positive otherwise. While the minimum value that A can get is when $p_2^r = 1$. Consequently $A^{min} = h$ and the first order at A^{min}, becomes

$$\frac{2}{3\left(1-\delta_1^2\right)}\left(\frac{h^4 - h^3\left(3+h\right)}{h^4}\right) < 0$$

If δ_1 is such that $3(1+h)\delta_1^4 > h(1+2\delta_1^3)$, the function is decreasing in p_2^r and the optimal p_2^{r*} is the corner solution $1 - \frac{1-\delta_1}{\delta_1}h$. Otherwise, p_2^{r*} is the interior solution, which is the first root of $h^3(1+2p_2^r+h) = (1-p_2^r+h)^4$. Therefore, for any δ_1, there exists a unique optimal p_2^{r*}. Given the expected revenue from sending a risky product, the firm sends a risky product if and only if $\pi_2\left(p_2^{r*}\left(\delta_1\right)\right) > \frac{1}{2}$.

Finally, we consider the second period's subgame where the firm sends a risky product in the first period and the customer returns it. Given δ_1, upon observing that the customer returned a risky product in the first period, the firm's posterior pdf of δ, distributed on the interval $[\delta_1, 1]$, is:

$$f\left(\delta \mid \text{return}\right) = \frac{p\left(v=0 \mid \delta\right) f(\delta)\, d\delta}{\int_{\delta_1}^{1}\left(1-\delta\right) f(\delta)\, d\delta} = \frac{1-\delta}{1-\delta_1-\frac{1}{2}\left(1-\delta_1^2\right)}$$

The firm's expected revenue from sending a risky product upon observing return is given by:

$$\pi_2^r = \int\limits_{\delta_2}^{1} s p_2^r f\left(s\right) ds \;\; = \;\; \left(\int\limits_{\delta_2}^{1} \frac{s-s^2}{1-\delta_1-\frac{1}{2}\left(1-\delta_1^2\right)} ds\right) p_2^r \tag{2.8}$$

where $\delta_2 = \max\left(\frac{h}{1-p_2^r+h}, \delta_1\right)$. If $\frac{h}{1-p_2^r+h} > \delta_1 \rightarrow p_2^r > 1 - \frac{1-\delta_1}{\delta_1} h$, the firm maximizes:

$$\max_{p_2^r} \left(\frac{\frac{1}{6} - \frac{h^2}{2(1-p_2^r+h)^2} + \frac{h^3}{3(1-p_2^r+h)^3}}{1-\delta_1-\frac{1}{2}\left(1-\delta_1^2\right)}\right) p_2^r \tag{2.9}$$

The first order condition gives:

$$\frac{d\pi_2^r}{dp_2^r} \;=\; \frac{\frac{1}{6} + \frac{h^3}{3(1-p_2^r+h)^3} - \frac{h^2}{2(1-p_2^r+h)^2} + p_2^r\left(\frac{h^3}{(1-p_2^r+h)^4} - \frac{h^2}{(1-p_2^r+h)^3}\right)}{1-\delta_1-\frac{1}{2}\left(1-\delta_1^2\right)}$$

Substituting $1-p_2^r+h$ with A, the first order condition becomes:

$$\frac{d\pi_2^r}{dp_2^r} \;=\; \frac{A^4 + 3h^2A^2 - 2h^2(3+5h)A + 6h^3(1+h)}{6A^4\left(1-\delta_1-\frac{1}{2}\left(1-\delta_1^2\right)\right)}$$

Note $1-\delta_1-\frac{1}{2}\left(1-\delta_1^2\right) > 0$, therefore, to determine the sign of FOC, we focus on the sign of the numerator. Furthermore, we have $h < A < \frac{h}{\delta_1}$. If $6(1+h)\delta_1^3 < h(1+\delta_1) + 4h\delta_1^2$, then the first-order has a root at $A^4 + 3h^2A^2 - 2h^2(3+5h)A + 6h^3(1+h) = 0$. However, if $6(1+h)\delta_1^3 \geq h(1+\delta_1) + 4h\delta_1^2$, the revenue function is strictly increasing in A (decreasing in p_2^r).

To check for the optimally of A' at $A'^4 + 3h^2A'^2 - 2h^2(3+5h)A' + 6h^3(1+h) = 0$, when $6(1+h)\delta_1^3 < h(1+\delta_1) + 4h\delta_1^2$, we write the SOC as follows::

$$\frac{d^2\pi_2^r}{dp_2^{r^2}} = \frac{\frac{2h^3}{(1-p_2^r+h)^4} - \frac{2h^2}{(1-p_2^r+h)^3} + p_2^r\left(\frac{4h^3}{(1-p_2^r+h)^5} - \frac{3h^2}{(1-p_2^r+h)^4}\right)}{1-\delta_1-\frac{1}{2}\left(1-\delta_1^2\right)}$$

Substituting $1 - p_2^r + h$ with A, the numerator becomes:

$$\frac{2h^3}{A^4} - \frac{2h^2}{A^3} + (1 + h - A)\left(\frac{4h^3}{A^5} - \frac{3h^2}{A^4}\right) = h^2\frac{A^2 - A(3+5h) + 4h(1+h)}{A^5}$$

When $6(1+h)\delta_1^3 < h(1+\delta_1) + 4h\delta_1^2$, at A', the SOC has a negative sign and the function is maximized at A'.

Therefore, for any value of δ_1, there exists a unique p_2^{r*} such that the firm sends a risky product if $\pi_2\left(p_2^{r*}\left(\delta_1\right)\right) > \frac{1}{2}$. The customer resumes the service if her $\delta > \frac{h}{1-p_2^{r*}+h}$, keeps the risky product if her realized value is 1 and returns it otherwise. $\square$

Proof of Proposition 2.1.

Proof. We write the firm's expected revenue from sending a risky product in period 1 as follows:

$$\pi_1^r(p_1^r) \quad = \quad \int_{\delta_1}^{1} \left(p_1^r + \pi_2\left(\delta \mid \text{keep}\right)\right) s\, ds + \int_{\delta_1}^{1} \pi_2\left(f\left(\delta \mid \text{return}\right)\right)(1-s)\, ds$$

Note, as described in the proof of 2.2, in any induced period 2 subgame, $p_2^{r*}\left(\delta_1\right)$ is such that $\delta_2^*\left(\delta_1\right) \geq \delta_1$. Therefore, a customer of type δ_1, either resume the service and gets an expected utility equal to 0, or leaves the service. When $h < p_1^r \leq 1$, the customer of type δ_1 expects $U_1(\delta_1) = \delta_1(1-p_1^r) - (1-\delta_1)h$ from joining the service, which gives: $\delta_1 = \frac{h}{1-p_1^r+h}$.

To analyze the subgame induced in period 1, we must first consider the subgame induced in period 2, given $\delta_1 = \frac{h}{1-p_1^r+h}$.

As such:

- Upon observing *"keep"* in period 1, from Lemma 2.2 we have: when $3(1+h)\delta_1^4 < h(1+2\delta_1^3)$, p_2^{r*} is the first root of $h^3(1+2p_2^r+h) = (1-p_2^r+h)^4$. The firm sends a risky product in the second period if the corresponding revenue is greater than $\frac{1}{2}$. Plugging in the value of δ_1 and the induced p_2^{r*}, we have: $3(1+h)\delta_1^4 < h(1+2\delta_1^3) \rightarrow 3h^3(1-2p_1^r)+6h^2(1-p_1^r)^2+4h(1-p_1^r)^3+(1-p_1^r)^4 > 0$ and the expected revenue function of the period 2 subgame (given in equations

2.5 and 2.6) is given by:

$$\pi_2\left(p_1^r\right) = p_2^{r*} \frac{2\left(1 - \frac{h^3}{\left(1-p_2^{r*}+h\right)^3}\right)}{3\left(1 - \frac{h^2}{\left(1-p_1^r+h\right)^2}\right)} \tag{2.10}$$

The firm sends a risky product if $\pi_2\left(p_1^r\right) > \frac{1}{2}$. When $3h^3(1-2p_1^r) + 6h^2(1-p_1^r)^2 + 4h(1-p_1^r)^3 + (1-p_1^r)^4 < 0$, $p_2^{r*} = 1 - \frac{1-\delta_1}{\delta_1}h$; plugging in $\delta_1 = \frac{h}{1-p_1^r+h}$, we have $p_2^{r*} = p_1^r$. The expected revenue from sending a risky product is given by:

$$\pi_r^2(p_1^r) = \left(p_1^r \frac{2\left(1 - \frac{h^3}{\left(1-p_1^r+h\right)^3}\right)}{3\left(1 - \frac{h^2}{\left(1-p_1^r+h\right)^2}\right)}\right) = \frac{2}{3}p_1^r\left(1 - \frac{h}{1-p_1^r+2h} + \frac{h}{1-p_1^r+h}\right) \tag{2.11}$$

The firm sends a risky product if $\frac{2}{3}p_1^r\left(1 - \frac{h}{1-p_1^r+2h} + \frac{h}{1-p_1^r+h}\right) > \frac{1}{2}$.

- Upon observing *"return"* in period 1, from Lemma 2.2 we have: if $6(1+h)\delta_1^3 < h(1+\delta_1) + 4h\delta_1^2$, then, p_2^{r*} is the first root of $3h^2(3p_2^r - 1) - 4h(1-p_2^r)^2 - (1-p_2^r)^3 = 0$. Replacing δ_1 with $\frac{h}{1-p_1^r+h}$, we have $6(1+h)\delta_1^3 < h(1+\delta_1) + 4h\delta_1^2 \rightarrow 3h^2(3p_1^r - 1) - 4h(1-p_1^r)^2 - (1-p_1^r)^3 < 0$. The firm sends a risky product in the second period if the corresponding revenue is greater than $\frac{1}{2}$. Plugging in the value of δ_1 and the induced p_2^{r*}, in the expected revenue function of the period 2 subgame (given in equations 2.8 and 2.9), we have:

$$\left(\frac{\frac{1}{6} - \frac{h^2}{2(1-p_2^{r*}+h)^2} + \frac{h^3}{3(1-p_2^{r*}+h)^3}}{1 - \frac{h}{1-p_1^r+h} - \frac{1}{2}\left(1 - \frac{h^2}{(1-p_1^r+h)^2}\right)}\right) p_2^{r*} < \frac{1}{2}$$

Our numerical analysis shows for any p_1^r that satisfies $3h^2(3p_1^r - 1) - 4h(1-p_1^r)^2 - (1-p_1^r)^3 < 0$, the expected revenue from sending a risky product is less than $\frac{1}{2}$. Therefore, when p_1^r is such that $3h^2(3p_1^r - 1) - 4h(1-p_1^r)^2 - (1-p_1^r)^3 < 0$, following a return observation, the firm sends a safe product priced at $\frac{1}{2}$, in the second period.

Meanwhile, when $3h^2(3p_1^r - 1) - 4h(1-p_1^r)^2 - (1-p_1^r)^3 > 0$, upon a return observation $p_2^{r*} = p_1^r$. Plugging in δ_1 and p_2^{r*}, the expected revenue from sending a risky product (given in equations 2.8 and 2.9) is:

$$\left(\frac{\frac{1}{6}-\frac{h^2}{2(1-p_1^r+h)^2}+\frac{h^3}{3(1-p_1^r+h)^3}}{1-\frac{h}{1-p_1^r+h}-\frac{1}{2}\left(1-\frac{h^2}{(1-p_1^r+h)^2}\right)}\right)p_1^r=\frac{p_1^r(1-p_1^r+3h)}{3(1-p_1^r+h)} \tag{2.12}$$

The firm sends a risky product upon observing return if and only if $\frac{p_1^r(1-p_1^r+3h)}{3(1-p_1^r+h)}>\frac{1}{2}$.

Suppose p_1^r is such that:

- The firm sends a risky product upon a *keep* observation priced at p_2^{r*} where p_2^{r*} is the first root of $h^3(1+2p_2^r+h)=(1-p_2^r+h)^4$. Therefore, $\delta_1<\delta_2\left(p_2^{r*}\right)$ and $3(1+h)\delta_1^4<h(1+2\delta_1^3)$. Moreover, the expected revenue from sending a risky product in the second period, given by equation 2.10, must be greater than the expected revenue from sending a safe product (i.e. $\frac{1}{2}$).
- The firm sends a safe product priced at $\frac{1}{2}$ upon a *return* observation. Therefore, the expected revenue from sending a risky product in the second period, given by equation 2.12, is weakly smaller than the expected revenue from sending a safe product (i.e. $\frac{1}{2}$).

In the first period, the firm would wish to maximize the expected revenue function given by:

$$\max_{p_1^r}\int_{\delta_1}^{1}\left(p_1^r+\pi_2\left(\delta\mid\text{keep}\right)\right)s\,ds\;+\int_{\delta_1}^{1}\pi_2\left(f\left(\delta\mid\text{return}\right)\right)(1-s)\,ds$$

$$s.t.\;\;\pi_2\left(\delta\mid\text{keep}\right)>\frac{1}{2},$$

$$\pi_2\left(\delta\mid\text{return}\right)\leq\frac{1}{2}.$$

Given the equilibrium path described above we have, $\pi_2\left(\delta\mid\text{return}\right)=\frac{1}{2}$ and $\pi_2\left(\delta\mid\text{keep}\right)=\int_{\delta_2\left(p_2^{r*}\right)}^{1}\frac{2s^2}{1-\delta_1^2}p_2^{r*}$. We rewrite the firm's expected revenue (using equation

2.5) as follows:

$$\begin{aligned}
\pi_1^r(p_1^r) &= \int_{\delta_1}^{1} \left(p_1^r + \pi_2\left(\delta \mid \text{keep}\right)\right) s\, ds \; + \int_{\delta_1}^{1} \pi_2\left(f\left(\delta \mid \text{return}\right)\right)(1-s)\, ds \\
&= \int_{\delta_1}^{1} s p_1^r ds + \int_{\delta_2\left(p_2^{r*}\right)}^{1} s^2 p_2^{r*} ds + \int_{\delta_1}^{1} (1-s)\frac{1}{2} ds \\
&= \left(\frac{1}{2} - \frac{h^2}{2\left(1-p_1^r+h\right)^2}\right) p_1^r + \left(\frac{1}{3} - \frac{h^3}{3(1-p_2^{r*}+h)^3}\right) p_2^{r*} + \\
&\quad \left(\frac{1}{2} - \frac{h}{1-p_1^r+h} + \frac{h^2}{2\left(1-p_1^r+h\right)^2}\right)\frac{1}{2}
\end{aligned}$$

The first expression in the revenue function denotes the expected revenue from period 1 where with the joint probability $\int_{\delta_1}^{1} sds = \frac{1}{2} - \frac{h^2}{2\left(1-p_1^r+h\right)^2}$, the customer joins and keeps the product and the firm earns p_1^r. The second expression denotes the expected revenue from period 2 given that the customer keeps the product in period 1. Finally, the last expression denotes the expected revenue from sending a safe product priced at $\frac{1}{2}$ in period 2, if the customer returns the product in period 1, with the joint probability $\int_{\delta_1}^{1}(1-s)ds = \frac{1}{2} - \frac{h}{1-p_1^r+h} + \frac{h^2}{2\left(1-p_1^r+h\right)^2}$. Note p_2^{r*} is the root of $h^3(1+2p_2^{r*}+h) = (1-p_2^{r*}+h)^4$, and is independent of p_1^r. While the second-order condition gives $\frac{d^2\pi_1^r}{dp_1^{r2}} = -h\frac{4h^2+3h+2-2(1-h)p_1^r}{2(1-p_1^r+h)^4} < 0$. Therefore, the revenue function is concave in p_1^r. We write the FOC as follows:

$$\begin{aligned}
\frac{d\pi_1^r}{dp_1^r} &= \frac{1}{2} - \frac{h^2}{2(1-p_1^r+h)^2} - \frac{h^2 p_1^r}{(1-p_1^r+h)^3} + \frac{1}{2}\left(\frac{h^2}{(1-p_1^r+h)^3} - \frac{h}{(1-p_1^r+h)^2}\right) \\
&= \frac{-p_1^{r3} + 3(1+h)p_1^{r2} - (3+5h+4h^2)p_1^r + 1 + 2h + 2h^2}{2(1-p_1^r+h)^3}
\end{aligned}$$

Therefore, the expected revenue has a local maxima which satisfies $-p_1^{r3} + 3(1+h)p_1^{r2} - (3+5h+4h^2)p_1^r + 1 + 2h + 2h^2 = 0$. Note given p_1^{r*}, and p_2^{r*}, all the belief consistency constraints are met for any values of h.

In short, in the first period, the firm's optimal pricing strategy is to announce p_1^{r*}. The customer joins the service of her $\delta \geq \frac{h}{1-p_1^{r*}+h}$. The firm sends a risky product. Upon realizing her value, the customer keeps the product if her realized value is 1 and returns it otherwise. The period 2 subgame depends on the action of the customer in period one such that:

- Upon a *keep* observation, the firm announces p_2^{r*} where p_2^{r*} is the first root of $h^3(1+2p_2^r+h) = (1-p_2^r+h)^4$. The customer resumes the service if her $\delta \geq \frac{h}{1-p_2^{r*}+h}$. The firm sends a risky product, which the customer keeps if her realized value is 1 and returns otherwise.

- Upon a a *return* observation, the firm announces $p_1^{r*} = \frac{1}{2}$. The customer resumes the service. The firm sends a safe product, which the customer keeps.

Moreover, the firm's expected revenue as described in item is given by:

$$\begin{aligned} \pi_1^r(h) &= \left(\frac{1}{2} - \frac{h^2}{2\left(1-p_1^{r*}+h\right)^2}\right)p_1^{r*} + \left(\frac{1}{3} - \frac{h^3}{3(1-p_2^*+h)^3}\right)p_2^* + \\ &\quad \left(\frac{1}{2} - \frac{h}{1-p_1^{r*}+h} + \frac{h^2}{2\left(1-p_1^{r*}+h\right)^2}\right)\frac{1}{2} \end{aligned} \tag{2.13}$$

Note if in the first period the firm announces $p_1^r \leq h$, the customer joins if her $\delta \geq p_1^r$. The firm sends a risky product, but does not update its belief about the customer's taste. Moreover, since $h \leq \frac{1}{2}$, the expected revenue is strictly smaller than the one described above, ruling out the pricing strategy ($p_1^r \leq h$) as optimal. Finally, if the firm sets p_1^r such that the optimal period 2 subgame following a *keep* observation, is for the firm to send a safe product, then it is intuitive that the firm does not benefit from learning and the strategy is inferior to the optimal strategy described above. □

Proof of Proposition 2.2.

Proof. At end of period 1, if the firm sent a safe product and the customer kept it, the firm does not gain any new information about δ and believes it to be uniformly distributed on interval $[0,1]$. According to the proof of Lemma 2.2, in period 2, the firm sends a safe product priced at $\frac{1}{2}$, which the customer keeps. Therefore, the firm earns $\frac{1}{2}$ in the period 2 subgame. In period 1, given p_1^s, the customer joins if $p_1^s \leq \frac{1}{2}$. Therefore, $p_1^{s*} = \frac{1}{2}$. The customer joins and the firm sends a safe product in both periods. The firm earns $\pi_1^s = 1$ in revenue. □

Proof of Theorem 2.1:

Proof. In Proposition 2.2, that p_1^{s*} is unique and $p_1^{s*} = \frac{1}{2}$. Finally, the firm sends a risky product in the first period if and only if its expected revenue from sending a risky product (equation 2.13 is greater than that of sending a safe product in both periods (i.e. 1). As such, we have:

$$\pi_1^r(h) > \pi_1^s$$

Note that at $h = 0$, the firm's expected revenue from the risky product equals $\frac{13}{12}$. Using the envelope theorem, it immediately follows that the firm expected revenue, evaluated at p_1^{r*} and p_2^{r*}, decreases in h:

$$\frac{d\pi_1^r}{dh} = \frac{\partial \pi_1^r}{\partial p_1^{r*}}\frac{dp_1^{r*}}{dh} + \frac{\partial \pi_1^r}{\partial p_2^{r*}}\frac{dp_2^{r*}}{dh} + \frac{\partial \pi_1^r}{\partial h} = \frac{\partial \pi_1^r}{\partial h}$$

Note, the revenue function is concave and continuous at p_1^{r*} and p_2^{r*}. Therefore, we have: $\frac{\partial \pi_1^r}{\partial p_1^{r*}} = 0$ and $\frac{\partial \pi_1^r}{\partial p_2^{r*}} = 0$. Meanwhile, we have:

$$\frac{\partial \pi_1^r}{\partial h} = -\frac{(1-p_1^{r*})(1-(1-2h)p_1^{r*})}{2(1-p_1^{r*}+h)^3} - \frac{h^2(1-p_2^{r*})p_2^{r*}}{(1-p_2^{r*}+h)^4} < 0 \tag{2.14}$$

Therefore, there exists the threshold h^{dp} such that the firm sends a risky product in period 1 (exploration equilibrium) if $h < h^{dp}$ and a safe product (exploitation equilibrium), otherwise. □

Appendix B: Fixed Pricing

Proof of Lemma 2.3.

Proof. Suppose a customer joins the service if her $\delta \geq \delta_1$. If the firm sends a safe product in period 1, the customer keeps it, regardless of her δ. Thus, the firm's posterior belief of δ remains unchanged (i.e. $f(\delta) \sim U[\delta_1, 1]$). In the second period, expecting a safe product, the customer of type δ resumes the service if $p^s \leq \frac{1}{2}$, and keeps the product if she receives it. Given $f(\delta)$, the firm's expected revenue from sending a safe product equals $\pi_2^s = p^s$.

Let us consider the subgame where given the updated belief $f(\delta)$, the firm sends a risky product. The customer keeps a risky product if her realized value is 1 and returns it otherwise. Moreover, the customer resumes the service if $\delta(1-p^r) - (1-\delta)h \geq 0$. Therefore, the firm's expected revenue from sending a risky product equals $\pi_2^r = \int_{\delta_2}^1 p^r f(s)ds$, where, $\delta_2 = \max\left(\frac{h}{1-p^r+h}, \delta_1\right)$. The firm sends a risky product in period

2, if and only if the expected revenue from sending a risky product is greater than that of sending a safe product (i.e. $\pi_2^r(\delta_1, p^r) > p^s$).

If the firm sends a risky product in the first period and the customer keeps it, the posterior pdf is $f(\cdot) = \frac{2\delta}{1-\delta_1^2}$. Given $f(\delta)$, the firm's expected revenue from sending another risky product in the first period is given by:

$$\pi_2^r = \int_{\delta_2}^{1} sp^r f(s)\,ds \quad = \quad p^r \int_{\delta_2}^{1} \frac{2s^2}{1-\delta_1^2} ds = p^r \frac{2\left(1-\delta_2^3\right)}{3(1-\delta_1^2)} \tag{2.15}$$

where $\delta_2 = \max\left(\frac{h}{1-p^r+h}, \delta_1\right)$. Therefore, the firm sends a risky product if and only if the expected revenue from sending a risky product is greater than that of a safe product.

Meanwhile, if the firm sends a risky product in period 1 and the customer returns it, the firm's posterior belief has the following distribution: $f(\delta) = \frac{1-\delta}{1-\delta_1-\frac{1}{2}\left(1-\delta_1^2\right)}$. Given $f(\delta)$, the firm's expected revenue from sending another risky product in the second period after the customer returned one in the first period is given by:

$$\begin{aligned}\pi_2^r &= \int_{\delta_2}^{1} sp^r f(s)\,ds = p^r \int_{\delta_2}^{1} \frac{(1-s)s}{1-\delta_1-\frac{1}{2}\left(1-\delta_1^2\right)} ds \\ &= p^r \frac{(1-\delta_2)^2(1+2\delta_2)}{3(1-\delta_1)^2} \end{aligned} \tag{2.16}$$

where $\delta_2 = \max\left(\frac{h}{1-p^r+h}, \delta_1\right)$. Again, the firm sends a risky product if and only if the expected revenue is greater than p^s. ☐

Proof of Proposition 2.3:

Proof. Suppose the firm sends a risky product priced at p^r. To derive δ_1 (i.e. the δ above which the customer joins the service in period 1), we derive the customer's expected utility given the equilibrium path:

Suppose p^r is such that in period 2, the firm sends a risky product upon observing *keep* in period 1, and a safe product upon observing *return*. Based on Lemma 2.3, upon keeping the product, the customer resumes the service if her $\delta \geq \delta_2$, where, $\delta_2 = \max\left(\frac{h}{1-p^r+h}, \delta_1\right)$. Meanwhile, upon returning the safe product in period 1, she resumes the service if $p_s \leq \frac{1}{2}$. Therefore, we have,

- If $\delta \in [0, \delta_2]$ then the customer's expected utility from joining the service and receiving a risky product in period 1 is $\delta\left(1-p^r\right)+(1-\delta)\left(\frac{1}{2}-p^s-h\right)$.
- If $\delta \in [\delta_2, 1]$ then the customer's expected utility is $\delta\left(1-p^r\right)-(1-\delta)\,h+\delta\left(\delta\left(1-p^r\right)-(1-\delta)\,h\right)+(1-\delta)\left(\frac{1}{2}-p^s\right) \geq 0$. Therefore, a customer of type δ where $\delta \geq \delta_2$ expects a non-negative utility from joining the service.

If $p^s > \frac{1}{2}-h$, $\delta_1 = \frac{p^s-\frac{1}{2}+h}{\frac{1}{2}-p^r+p^s+h}$ and if $p^s \leq \frac{1}{2}-h$, $\delta_1 = 0$. Regardless, we have $\frac{h}{1-p^r+h} > \delta_1$, therefore, $\delta_2 = \frac{h}{1-p^r+h}$. The firm's expected revenue from sending a risky product in period 1 is given by:

$$\pi_1(r) \;=\; \int\limits_{\delta_1}^{1} \left(p^r+\pi_2\left(\delta \mid \text{keep}\right)\right) s\, ds \;+ \int\limits_{\delta_1}^{1} \pi_2\left(f\left(\delta \mid \text{return}\right)\right)(1-s)\, ds$$

where, from equations 2.15 and 2.16, we have $\pi_2(keep) = p^r\frac{2\left(1-\delta_2^3\right)}{3(1-\delta_1^2)}$ and $\pi_2(return) = p^r\frac{(1-\delta_2)^2(1+2\delta_2)}{3(1-\delta_1)^2}$. Plugging in the values of $\pi_2(keep)$ and $\pi_2(return)$, the firm maximization problem becomes:

$$\max_{p^r,p^s}\left(\frac{1}{2}\left(1-\delta_1^2\right)p^r+\frac{1}{3}\left(1-\frac{h^3}{(1-p^r+h)^3}\right)p^r+\left(\frac{1}{2}-\delta_1+\frac{1}{2}\delta_1^2\right)p^s\right)$$

$$s.t.\;\; p^r\frac{2\left(1-\frac{h^3}{(1-p^r+h)^3}\right)}{3(1-\delta_1^2)} > p^s$$

$$p^r\frac{(1-\frac{h}{1-p^r+h})^2(1+2\frac{h}{1-p^r+h})}{3(1-\delta_1)^2} \leq p^s$$

Suppose, the firm sets $p^s \geq \frac{1}{2}-h$, and $\delta_1 = \frac{p^s-\frac{1}{2}+h}{\frac{1}{2}-p^r+p^s+h}$. When, h is less than h', where h' is the second root of $-5+20h'-30h'^2+22h'^3+5h'^4=0$, the expected revenue is strictly decreasing in p^s. Thus, $p^{s*} = \frac{1}{2}-h$, $\delta_1 = 0$ and the revenue function is concave in p^r and is maximized at p^{r*}, where p^{r*} is the first root of $5(1-p^r+h)^4 = 2h^3(1+2p^r+h)$. However, when $h > h'$, the revenue function has a unique p^{r*} and p^{s*}, where: p^{r*} and p^{s*} jointly solve equations 2.17 and 2.18.

$$p^{r*}(1+2p^{s*}+6h)-2h-2p^{r*2}+2p^{s*}-1=0 \qquad (2.17)$$

$$\frac{1}{3}-\frac{h^3(1+2p^{r*}+h)}{3(1-p^{r*}+h)^4}-\frac{2p^{r*}}{1-2p^{r*}+2p^{s*}+h}+\frac{2(1-p^{r*})(2+p^{r*}+2h)}{(1-2p^{r*}+2p^{s*}+h)^2}-\frac{4(1+2h)(1-p^{r*})^2}{(1-2p^{r*}+2p^{s*}+h)^3}=0 \qquad (2.18)$$

Next we check whether p^{r*} and p^{s*} meet the constraints. When $h \leq h'$, ($p^{s*} = \frac{1}{2}-h$

and p^{r*} is the first root of $5(1-p^r+h)^4 = 2h^3(1+2p^r+h))$, the numerical analysis shows for any $h \leq h'$, p^{s*} and p^{r*} satisfy both constraints. However, when $h > h'$, the second constraint is only satisfied when $h <$ 0.464205. Therefore, for $h >\approx 0.464205$, our numerical analysis indicates that p^{s*} binds the second constraint.

In short, for $h \leq h'$, where $h' \approx 0.43336$ (which is of interest to us), $p^{s*} = \frac{1}{2} - h$ and p^{r*} is the first root of $5(1-p^r+h)^4 = 2h^3(1+2p^r+h)$. The firm sends a risky product priced at p^{r*} in the first period. The customer joins the service regardless of her δ. The firm sends a risky product. The customer keeps the product if her realized value is 1 and returns it otherwise. The firm sends a risky product upon observing *keep*, and a safe product otherwise. The firm's expected revenue is given by:

$$\pi_1^r = \frac{1}{2}p^{r*} + \frac{1}{3}\left(1 - \frac{h^3}{(1-p^{r*}+h)^3}\right)p^{r*} + \frac{1}{2}p^{s*} \tag{2.19}$$

□

Proof of Proposition 2.4:

Proof. When the firm sends a safe product in the first period, its belief about the customer's taste remains unchanged and therefore, as described in Lemma 2.3, the firm sends another safe product in the second period. The optimal price in such a case is $p^{s*} = \frac{1}{2}$ where the firm sends a safe product in both periods and extracts all surplus. Therefore the optimal revenue is 1. This subgame equilibrium is the unique period 1's safe product subgame equilibrium. □

Proof of Theorem 2.2:

Proof. Given the results of Proposition 2.3 and Proposition 2.4, the firm sends a risky product in the first period if and only if its expected revenue from sending a risky product (equation 2.19) is greater than that of sending a safe product in both periods (i.e. 1). As such, we have:

$$\pi_1^r(h) > \pi_1^s$$

Note that at $h = 0$, the firm's expected revenue from the risky product equals $\frac{13}{12}$ and $\frac{d\pi_2^r}{dh} < 0$. Using the envelope theorem, it immediately follows that the firm expected revenue, evaluated at p^{s*} and p^{r*}, decreases in h:

$$\frac{d\pi_1^r}{dh} = \frac{\partial \pi_1^r}{\partial p^{r*}}\frac{dp^{r*}}{dh} + \frac{\partial \pi_1^r}{\partial p^{s*}}\frac{dp^{s*}}{dh} + \frac{\partial \pi_1^r}{\partial h} = \frac{\partial \pi_1^r}{\partial h}$$

Note, the revenue function is concave and continuous at p^{r*} and p^{s*}. Therefore, we have: $\frac{\partial \pi_1^r}{\partial p^{r*}} = 0$ and $\frac{\partial \pi_1^r}{\partial p^{s*}} = 0$. Meanwhile, we have:

$$\frac{\partial \pi_1^r}{\partial h} = -(1 - p^{r*})\left(\frac{h^2 p^{r*}}{(1 - p^{r*} + h)^4} + \frac{4(2p^{s*} - (1 - 2h)p^{r*})}{(1 - 2p^{r*} + 2p^{s*} + 2h)^3}\right) < 0 \tag{2.20}$$

Therefore, there exists the threshold h^{fp} such that the firm sends a risky product in period 1 (exploration equilibrium) if $h < h^{fp}$ and a safe product (exploitation equilibrium) otherwise. □

Appendix C

Proof of Proposition 2.5:

Proof. Note at $h = 0$, the optimal expected revenue under dynamic and fixed pricing equals $\frac{13}{12}$. Given equations 2.14 and 2.21, we have for any h:

$$\frac{\partial \pi_1^{fp}}{\partial h} > \frac{\partial \pi_1^{dp}}{\partial h} \tag{2.21}$$

Therefore, $\pi_1^{fp} > \pi_1^{dp}$, and $h^{fp} > h^{dp}$. □

Proof of Proposition 2.6:

Proof. Under a no return policy, suppose a customer joins the service if her $\delta \geq \delta_1$. Given that in the first period, the customer has no other option than to keep the product, the firm's posterior belief of δ remains unchanged (i.e. $f(\delta) \sim U[\delta_1, 1]$). In the second period's subgame, if the firm announces p_2^r and is expected to send a risky product, the customer resumes the service if her $\delta \geq p_2^r$. Therefore, the firm's expected revenue equals $\pi_2^r = \int_{\delta_2}^{1} p_2^r$, where $\delta_2 = \max(\delta_1, p_2^r)$. $\pi_2^r \leq \frac{1}{2}$ for any δ_2. Consequently, it is never optimal for the firm to send a risky product in the second period. Similar results are achieved in the first period's subgame. Therefore, under a no return policy, the firm's optimal action is to send a safe product in both periods. □

Proof of Lemma 2.4:

Proof. Under dynamic pricing with α, in the second period subgame, the customer resumes the service if $\delta \geq \delta_2$, where $\delta_2 = \frac{(1-\alpha)h}{1-p_2^r+(1-\alpha)h}$. Given the posterior belief distribution $f(\cdot)$, the firm's optimal price for the risky product maximizes:

$$\left(\int_{\delta_2}^{1} sf(s)ds \right) (p_2^r + \alpha h) - \alpha h \tag{2.22}$$

With α, under dynamic pricing, the proofs for the existence and the uniqueness of p_2^{r*}, in spirit, are similar to the proofs provided for Lemma 2.2. To solve for the first period subgame, we analyze the possible induced subgames given the customer's behavior in period 1. Let us consider the first period subgame where the firm sends a risky product. The customer keeps a risky product if her realized value is 1 and returns it otherwise. If the firm sent a risky product in the first period and the customer kept it, the firm sends a risky product priced at p_2^{r*} which is the first root of polynomial $(1 - p_2^r + (1-\alpha)h)^4 = (1-\alpha)^3 h^3 (1 - 4p_2^r + (1-4\alpha)h)$. The firm's expected revenue from sending a risky product is given by:

$$\pi_1^r(p_1^r) = \int_{\delta_1}^{1} \left(p_1^r + \pi_2\left(\delta \mid \text{keep}\right)\right) s \, ds \; + \int_{\delta_1}^{1} \left(-\alpha h + \pi_2\left(f\left(\delta \mid \text{return}\right)\right)\right)(1-s)\, ds$$

Where $\pi_2\big(f\left(\delta \mid \text{keep}\right)\big) = (p_2^{r*} + \alpha h) \int_{\delta_2(p_2^{r*})}^{1} sf(s)\, ds - \alpha h$, $f\left(\delta \mid \text{keep}\right) = \frac{2\delta}{1-\delta_1^2}$ and $\pi_2\big(f\left(\delta \mid \text{return}\right)\big) = \frac{1}{2}$. Plugging in the values, we have that in the first period subgame, p_1^{r*} is the solution to the maximization problem given by:

$$\max_{p_1^r} \left(\frac{1}{2}\left(1-\delta_1^2\right) \left(p_1^r + \frac{2}{3}\frac{\left(1-\delta_2^3\right)}{1-\delta_1^2}\left(p_2^{r*}+\alpha h\right) - \alpha h \right) + \left(\frac{1}{2} - \delta_1 + \frac{1}{2}\delta^2 \right) \left(-\alpha h + \frac{1}{2} \right) \right)$$

$$s.t. \;\; \frac{2}{3}\frac{\left(1-\delta_2^3\right)}{1-\delta_1^2}\left(p_2^{r*}+\alpha h\right) - \alpha h > \frac{1}{2}$$

$$\frac{\frac{1}{2}\left(1-\delta_1^2\right) - \frac{1}{3}\left(1-\delta_1^3\right)}{1-\delta_1-\frac{1}{2}\left(1-\delta_1^2\right)}\left(p_1^r+\alpha h\right) - \alpha h \leq \frac{1}{2}$$

$$h < p_1 \leq p_2^*$$

where $\delta_1 = \frac{(1-\alpha)h}{1-p_1^r+(1-\alpha)h}$ and $\delta_2 = \frac{(1-\alpha)h}{1-p_2^{r*}+(1-\alpha)h}$. Note p_2^{r*} does not depend on p_1^r. There exists a unique solution to the maximization problem at the second root of the polynomial of degree three: ${p_1^r}^3 - 3(1+(1-\alpha)h){p_1^r}^2 + (3+5(1-\alpha)h+2(1-\alpha)(2-\alpha)h^2)p_1^r - 2(1-\alpha)^2h^3 - 2(1-\alpha)h^2 - 2(1-\alpha)h - 1 = 0$ which satisfies the equilibrium constraints.

Using the envelope theorem, it follows that the firm expected revenue, evaluated at p_1^{r*} and p_2^{r*} is strictly increasing in α,

$$\frac{d\pi_1^r}{d\alpha} = \frac{\partial\pi_1^r}{\partial p_1^{r*}}\frac{dp_1^{r*}}{d\alpha} + \frac{\partial\pi_1^r}{\partial p_2^{r*}}\frac{dp_2^{r*}}{d\alpha} + \frac{\partial\pi_1^r}{\partial\alpha} = \frac{\partial\pi_1^r}{\partial\alpha} > 0$$

Note, the revenue function is concave and continuous at p_1^{r*} and p_2^{r*}. Therefore, we have: $\frac{\partial\pi_1^r}{\partial p_1^{r*}} = 0$ and $\frac{\partial\pi_1^r}{\partial p_2^{r*}} = 0$. Therefore, $\alpha^* = 0$.

Similar results are attained solving for the fixed-pricing equilibrium. $\square$

Chapter 3

Coordination of Streaming Platforms in the Presence of First-Party Content

1 Introduction

A video streaming platform operates as a two-sided market. The platform aggregates a bundle of content – either produced by the platform itself or sourced by third-party providers – and offers users unlimited access to the bundle in exchange for a subscription fee. Offering users with heterogeneous taste a large bundle of independent content has two advantages: first, users are more likely to find content that matches their taste; second, users who are uncertain about their valuation for each individual content, can predict their valuation for the collection of bundle more easily (i.e. with higher certainty). In short, this business model increases the platform's ability to extract consumer surplus by means of reducing the heterogeneity in users valuations (see Abdallah et al. (2021) for a recent review of the literature on bundling). However, as opposed to a typical two-sided market, in which one (or both) side pays a fee each time that a transaction occurs, in a subscription-based market, users pay a fixed price to enter the market and consume content at their own discretion (see Sanchez-Cartas and Leon (2019) for a comprehensive review of the literature on the pricing theory in multi-sided markets). Similar to a typical two-sided market, in the presence of strong indirect network effects, each individual provider benefits from the existence of other providers, which, collectively, incentivizes more users to join the service by virtue of its higher value. However, since the consumers pay for access to the complete bundle – and not an individual content – it is difficult to isolate the standalone revenue generated by the presence of a single provider.

In practice, streaming platforms employ a variety of revenue allocation rules to compensate providers. "*Revenue-sharing*" and "*licensing*" payment structures are the two most common. Under a *revenue-sharing* contract, the platform allocates a share of the total revenue generated from subscription fees to the third-party providers. This share is further divided amongst individual providers, proportional to the realized consumption of their content on the platform. The revenue sharing model is used by streaming platforms such as Amazon Prime, Spotify and Hulu. Under a *licensing* contract, the platform offers a lump-sum to secure the right to the content of a provider for an agreed-upon period of time[1]. The licensing revenue allocation rule is employed by streaming platforms such as Netflix and HBO MAX.

Not all content create the same externalities. That is a provider's content may contribute more to the utility function of the subscribers, but the provider may not necessarily be compensated for it. The presence of popular TV-shows or classic movies (hereafter denoted by "popular content") generates greater externalities than their "niche" counterparts. Consequently, providers of popular content demand higher shares of revenue from the platform in return for the exclusive rights to their content[2]. Therefore, to maintain their revenue, platforms offer heterogeneous royalty fees to providers[3]. Over the last few years, most streaming platforms have invested

---

[1] See Brian Beer, How Netflix Pays for Movie and TV Show Licensing on Investopedia.

[2] In 2019, Netflix lost the streaming rights to The Office as the producers of the show "Universal Television" made a "*take it or leave it*" $100 million per-year offer (for five years).

[3] Hulu's royalty rate for big content providers yield a 20%-30% of generated revenue to Hulu, but with smaller content providers the revenue split goes as high as 50%-50% https://www.cbsnews.com/news/hulu-makes-a-profit-video-content-owners-not-so-much/. Amazon Prime, on the other hand, pays content providers between 6 cents and 15 cents per hour of streaming. Specifically, for content that are streamed less than 10,000 hours the royalty rate is 6 cents; and for content that are streamed over 1 million hours, the rate is 15 cents: https://digiday.com/future-of-tv/amazon-royalties-video-makers-uploading-prime-video/

in creating their own content (hereafter denoted as first-party content) in response to the high fees associated with securing the rights to stream popular content. Therefore, their business models have transformed from being a third-party aggregator, to a hybrid model operating as a third-party aggregator and first-party content producer. This shift further complicates the coordination problem of the two-sided platform. As an experience good, whose quality cannot be ascertained by users before purchase, the production of media content exhibits significant uncertainty. Therefore, the platform faces a trade-off between obtaining third-party content at a high cost, or investing in the production of her own uncertain content.

The aim of the present paper is thus twofold. First, we explore the performance of prevalent revenue allocation rules (*revenue-sharing* and *licensing*) in coordinating the two-sided platform[4]. Second, we investigate when it is optimal for the platform to invest in the risky production of first-party content, and the investment's impact on the platform's optimal content-mix, prices, and welfare. To our knowledge, little is known in the nascent literature about the optimal coordination strategies of video streaming platforms, and their interplay with the platform's strategic investment in the risky production of first-party content. There are three fundamental building blocks that distinguish our paper from the literature on coordination of two-sided markets, described as follows:

- Sequential entry: theoretical research mostly assumes that the two sides enter the market simultaneously, and derives the market equilibrium, accordingly

---

[4] Two-sided platforms must coordinate the interdependent joining decisions of the two sides of the market. In a simultaneous game, a coordinating pricing strategy solves the following chicken & egg problem: suppose the stand-alone utility of users on both sides from joining the platform is negative. To induce buyers (sellers) to join, the platform must convince the sellers (buyers) join as well. The market faces unfavorable expectation if given any pricing strategy no user joins from either side. The literature on two-sided market studies this problem and suggests strategies to overcome the unfavorable expectation (Caillaud and Jullien, 2003). In our problem, however, given the sequential entry of the two-sides, a coordinating contract between the platform and the providers maximizes the total revenue generated from the bundle in the resulting subgame perfect equilibrium.

(Parker and Van Alstyne, 2005; Weyl, 2010)[5]. In our setting however, given that a user must pay the subscription fee *ex ante*, and consume the content *ex post*, simultaneous entry does not apply. Specifically, the platform must first attract the providers to join; and only once the content-mix is determined can the platform present the user with the bundle to attract her.

- Revenue structure: in the prior two-sided markets literature, sellers set their own prices and derive revenue directly from each interaction with their buyer; and the platform's revenue is determined from the royalty fee collected from each transaction. However, in a bundling and subscription setting, providers relinquish their right to setting their own prices, and the platform governs the price for its entire bundle. Providers are paid separately once the platform collects all revenue generated from the user side.

- First-party content: research in the two-sided platforms literature that examines the platform's investment in first-party content is scarce. Close to our work is Hagiu and Spulber (2013) which shows that the platform may overcome an undesirable market Nash equilibrium where *no buyers or sellers join* by attracting buyers through the strategic introduction of first-party content. Our paper is a sequential entry market, so 1) a *nobody joins* equilibrium does not exist; 2) we study the platform's strategic investment in the production of first-party content as a means to increase participation on both the user and provider side.

To achieve our goals, we develop a two-period game-theoretic model of a monopoly two-sided streaming platform that aggregates content from a single popular content

---

[5] Hagiu and Jullien (2011) assume a market where buyers enter the market after sellers (sequential entry) and study the incentives of an information intermediary to divert the search of buyers. Hagiu and Spulber (2013) study the platform's pricing strategies under sequential entry when the platform may offer some integrated content to appear more attractive to buyers.

provider and a mass of niche providers. In the first period, the platform makes a take-it-or-leave-it offer to the providers, which the providers can accept or reject, given their outside option. The outside option for niche providers is negligible, which allows the platform to procure their content at a low cost. On the other hand, the popular content provider has a non-negligible outside option, and joins the platform only if her expected revenue from joining the platform is not less than her outside option. In the second period, the content available on the bundle is known under each revenue-allocation rule, and the platform sets a subscription fee. Users are uncertain about their valuation for each individual content apriori, and realize their value only upon sampling the content. Therefore, users make their decision to join the platform based on their expected valuation of the entire bundle and the subscription fee set by the platform. Upon joining, users sequentially sample the content available to find a single unit of content that matches their taste. If a user is matched with a content, she stops sampling and consumes that content. On the other hand, if she does not find a content that matches her taste after sampling all available content, she leaves without consuming any. Under the revenue-sharing contract, the payoff to each provider is derived from users' consumption of their content. This model allows us to gain insights regarding the platform's optimal pricing and content-mix strategies under each revenue allocation rule.

We show that the two-sided market is coordinated with a licensing contract. In our problem a coordinating contract between the platform and the popular provider maximizes the total revenue generated from the bundle in the resulting subgame perfect equilibrium. Under licensing, the platform's choice of subscription fee and consequently, the total net revenue generated from the bundle in the second stage, is decoupled from the platform's choice of payment to providers in the first stage. Our analysis further shows that under revenue-sharing, the platform strategically manipulates the popular content's consumption rate by setting a higher subscription fee such that users with an exceedingly unfavorable taste for niche content do not join. Effectively, since the platform announces the subscription fee only after the providers have joined, the platform's optimal action is to maximize its revenue by reducing the payment to the popular provider. We show that a price commitment strategy – where the platform credibly commits to the price it will charge the subscribers when trying to attract the providers to join the platform – under revenue-sharing is able to address this problem, and coordinate the market. The advantage of revenue-sharing with price commitment is not surprising. Without price commitment, the popular provider correctly anticipates the platform's pricing strategy in the second stage, and subsequently demands a higher royalty fee which, in turn, decreases the platform's profit. Conversely, with price commitment the platform gives up her right to manipulate the consumption rate of content via pricing, and thus convinces the popular content provider to join with a smaller royalty fee; this maximizes the total generated revenue from the bundle.

Next, we incorporate the platform's strategic production of first-party content in our model. We endogenize the probability of the first-party content successfully becoming a popular content by characterizing it as an increasing function of the platform's investment. The quality of the first-party content – whether it is ultimately popular – is realized only after the providers have made their joining decisions. First, we find that the presence of the popular content provider generates higher revenue for the platform. Specifically, the presence of a popular provider guarantees that some of the users with unfavorable taste for niche content join the platform, regardless of whether the first-party production is successful (i.e. popular) or unsuccessful (i.e. niche). That said, we also find that when the platform (correctly) anticipates the participation of the popular content provider, she invests significantly more in her first-party production. This is because a successful production accompanied by the presence of the popular provider results in much higher demand and higher total revenue. To reach this equilibrium, the platform is also willing to compensate a higher outside option for the popular provider. However, if the popular provider's outside option cannot be compensated, the platform still invests in the production of first-party content, albeit, significantly less.

2 Literature Review

Our work relates to the stream of the literature on two-sided market and the design of media platforms. The literature on two-sided platforms highlights the cross-side externalities that spur growth on both side of the market (Parker and Van Alstyne, 2005) and provides a theoretical framework for analyzing these markets (Rochet and Tirole, 2006; Armstrong and Wright, 2007). Jiang et al. (2017) provide a comprehensive review of literature on two-sided market pricing in operations management and pose the sequential entry as a main future research direction. Closely related to our work is the body of literature that examines a media platform's optimal pricing strategies on both sides of the market. Li et al. (2020) studies the interaction between ad-supported media platforms and users and the implications of three different pricing models for selling digital music: ownership, subscription and mixed pricing models. Using a stylized model, they find that the optimal pricing strategy depends on the platform's advertising revenue rate and consumers disutility from consuming advertisements. This study focuses on the interaction between users and the platform, where the revenue generated from the subscription is derived from ads. The paper also omits the interaction between the platform and the supply side. Lei and Swinney (2018) studies the decision of creators of digital goods to distribute their content a la carte or via subscription services, optimally. They model the selling decision of two content creators (low and high quality) via the platform. The creators can choose to sell their content a la carte and/or via subscription. The platform compensates the creators via revenue sharing contracts with homogeneous royalty fees. They find that the platform cannot induce *only* the high quality creator to sell via subscription which entails that the content offered via subscription is of lower quality compared to a la cart. A few studies have investigated the design of optimal pricing for ad-supported media platforms in the presence of user's advertisements disutility (DeValve and Pekeč, 2016; Armstrong and Weeds, 2007). In this paper, however, we consider the interaction between users, the platform, and the content providers and characterize the optimal design of the platform, when 1) each side enters the market, sequentially, and 2) the platform collects revenue generated from the bundle and distribute a portion of the revenue amongst providers. Our work is also related to the stream of literature on the coordination of a two-sided platform when the platform invests in producing first-party content. In particular, Anderson Jr et al. (2014) investigate the platform's choice to invest in new product development, in the video game consoles industry with network externalities. They build a stylized model to study the trade-off between investing in high platform performance versus lowering the investment to facilitate the third-party content development. Hagiu and Spulber (2013) introduce the investment in first-party content as a strategic variable for two-sided platforms to make participation more attractive to buyers, independently of the presence of the sellers. The paper considers homogeneous sellers, all of whom join the platform if the profit is non-negative or none joins otherwise. Thus, the platform's pricing on the seller side is not taken as the decision variable in the profit maximization problem.

Furthermore, in this study, we leverage intuition from the rich literature on the bundling of information goods (Bakos and Brynjolfsson, 1999; Geng et al., 2005;

Bhargava, 2021; Abdallah et al., 2021). Most of the work on bundling investigate the benefits of pure bundling as well as the platform's optimal pricing strategies. For instance, in their seminal paper, Bakos and Brynjolfsson (1999) finds that bundling very large numbers of information goods with i.i.d consumer valuations achieves greater profit and economic efficiency (all consumer surplus is extracted). They argue that this is in part because of the law of large numbers which allows users to predict their expected utility from a large bundle of content much easier than their expected utility from individual content. Geng et al. (2005) examines the optimality of bundling when a user's utility from consuming the goods declines with the number of information goods consumed. They show as long as the consumer values do not decrease too quickly, bundling is approximately optimal. In a study of the economic structure of co-production markets, Bhargava (2021) draws insights from the bundling literature to build an algebraic expression of bundle demand. Similar to Bhargava (2021), we draw from the literature to design a user's utility function such that the utility of a user *increases* in the number of content available on the platform, but at *diminishing rate*. However, in contrast to Bhargava (2021), we allow for heterogeneity of taste amongst users and thus, heterogeneity in the externalities created by each content. Note that in this setting as apposed to a classic two-sided market, the utility of users is non-linearly increasing in the number of sellers (providers) on the other side of the market. Furthermore, we make the assumption that the platform has access to and may invite a large enough collection of content such that users find at least one content that matches their heterogeneous taste, albeit the value they get from the consumption of the content may vary. This assumption is particularly relevant in the context of digital goods where platforms offer users access to thousands of content in their libraries (Alaei et al., 2019).

3 Model Description

We model a monopolist two-sided streaming platform that offers users with heterogeneous tastes an unlimited access to a bundle of digital goods, for a flat fee. The platform may be an *aggregator*, which bundles content only from third-party providers, or a *hybrid*, which produces first-party content, in addition to being an aggregator. The market is populated by a continuum of users of total mass normalized to 1 and a large pool of content providers. We assume a single subscription period, during which users may *sample* as many products as they like, but ultimately *consume* at most one content from the bundle.

The platform may compensate its providers using one of two revenue allocation contracts: *revenue-sharing* or *licensing*. Under the revenue-sharing structure, in the first stage the platform offers to distribute a share of its revenue amongst providers proportional to the streaming rate of their content on the platform[6]. Under the licensing structure, the provider allocates revenue according to a previously agreed-upon lump-sum payment.

We model the interaction between the platform, providers and users in a sequential game. We provide an overview of the model when the platform is an aggregator and hybrid, in the remainder of this section.

As shown in Figure 3·1, the interaction between the aggregator platform, providers, and users occurs over two stages. At the beginning of the first stage, the platform chooses a subset of providers to invite to join the platform, and the payment structure by which providers are compensated. Next, the invited providers decide whether to join (denoted as the provider's joining decision). At the beginning of the second stage, the bundle of providers on the platform are known, and the platform announces

---

[6] This revenue allocation strategy is prevalent amongst media streaming platforms. See Lei and Swinney (2018) and Alaei et al. (2019) for a similar revenue sharing structure in the study of digital goods subscription services.

the subscription fee to users. The users observe the subscription fee and the bundle of content, and decide whether to subscribe to the platform (denoted as the users' joining decision). Finally, the revenue of the bundle is realized based on the fraction of users who joined, and the platform compensates its providers according to the payment structure.

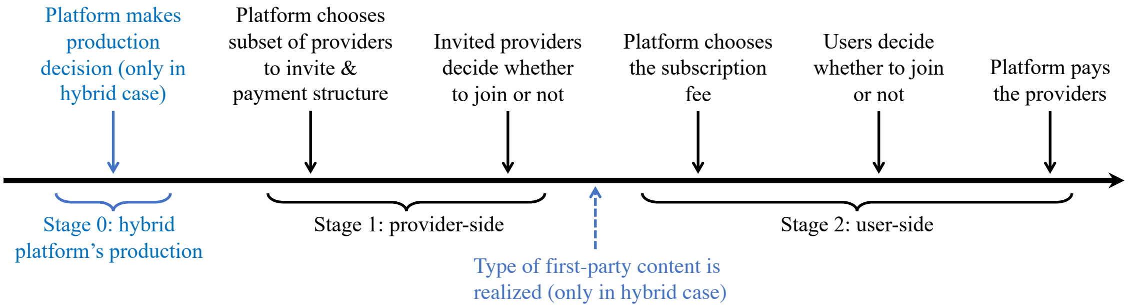

Figure 3·1: Timeline of events

We extend the baseline model to analyze the market equilibrium when a *hybrid* platform produces a unit of "first-party" content to add to its bundle. The development of first-party content is uncertain in that *apriori* the platform does not know the quality of the resulting production. As shown in Figure 3·1, the interaction between the hybrid platform, providers, and users occurs over three stages. In stage zero, the platform makes a production decision. In stage one, the platform chooses the subset of providers to invite to join the platform, and the payment structure. Subsequently, the providers who are invited, choose whether to join the platform. Prior to the second stage, the quality of the first-party production is realized. The game continues in the second stage in a similar manner as the baseline model.

3.1 Providers

For simplicity, we assume that each content provider offers one unit of content and therefore, the number of content available in the bundle equals the number of providers that join the platform[7]. We consider two types of third-party providers based on the content they offer: niche and popular. A niche provider owns content that caters to only a small set of users and therefore, has a high valuation dispersion amongst users. In contrast, a popular provider owns content that is “popular”, and caters to a broad range of users; therefore, this content has low valuation dispersion amongst users. For simplicity we assume there are a single popular provider and a sufficient pool of niche content providers, each of which may choose to participate in the platform[8]. A content provider faces a trade–off when making her content available on the platform. On one hand, she receives a payment from the platform to make her content available. On the other hand, she incurs an opportunity cost from offering her content through the platform and not selling it through other channels. We normalize the outside option of the niche providers to zero, but set the opportunity cost of the popular provider to θ.[9] This allows us to focus on the coordination problem of the platform when there is a large value creation disparity between a popular content compared to a niche content. Moreover, given that there exists a large number of niche providers, assuming a zero outside option for them is reasonable. The popular provider joins the platform if and only if her expected revenue is weakly greater than her outside option, θ. In section 6.1, we extend the model to allow for positive outside option for niche providers, and observe that our main results hold.

---

[7] Hagiu and Spulber (2013) and Hagiu and Wright (2015) make the same assumption when modeling a two-sided platform similar to our setting

[8] The assumption of large assortment of niche contents, is particularly valid in the context of digital goods and was first used in the seminal paper by (Bakos and Brynjolfsson, 1999) on the bundling of large number of information goods with zero marginal costs. As an example, there are more than 13,000 movies on Amazon Video Prime and the subscription service offers an estimate of 1,400 movies per dollar spent, while Netflix has over 15,000 titles in its catalog and offers 1,600 titles per dollar spent.

[9] Examples of such instances can be found in markets such as music or entertainment industry. For instance, popular artists such as Taylor Swift and Pink Floyd (Guardian, 2018) or producers of popular TV shows such as The Office and Friends (Times, 2018) have been able to negotiate for higher royalty fees as they generally appeal to a wider range of users than other lesser-known providers.

3.2 Users

A user realizes her value for a content only upon sampling it. A user expects to receive a value of 1 or 0 from the popular content, each with a probability of $\frac{1}{2}$. Meanwhile, she receives value x from a niche content with a probability of p, and 0 with a probability of $1 - p$. A user is said to be "matched" with a content if her realized value is greater than 0. We assume $p \ll \frac{1}{2}$ and that x is distributed across the population uniformly on the $(0, 2]$ interval. As such, a user of type x with $1 < x \leq 2$ receives a higher value (x) from consuming a matched niche content than a matched popular content. Therefore, she sequentially samples the niche contents available in the bundle to find a content that matches her taste (e.g browsing through contents in Netflix or reading movie descriptions); and if she does not find such content, she moves on to sample the popular content. On the other hand, a user of type x with $0 < x \leq 1$, receives a higher value from the popular content if it matches her taste and therefore, she samples the popular content first, and if she is not matched, she moves on to sequentially sampling from the niche contents available in the bundle to find one that matches her taste. A user's type x is known to her but is unknown to the platform.

3.3 The Motivation Behind The Users' Utility Function

Given the description above, an important question to address is *why wouldn't all consumers want to sample the popular content first?* Our model setup is motivated by the long tail phenomena. The "long tail" effect theorizes that an expanded product variety, increases the consumption of niche products, while decreasing the demand for popular products (Anderson, 2006; Brynjolfsson et al., 2011). The effect is most pronounced in industries such as the media industry, that generally have highly heterogeneous and variety-seeking consumers who enjoy different types of contents. The large size of the bundle of niche contents allows users to search for contents that are closest to their preferences and satisfy their heterogeneous tastes. This shifts away some consumers from a “hit” or “popluar” content towards niche contents. However, as empirically shown by Tan et al. (2017), not all consumers search for niche content, and some prefer to consume popular content. Hence, allowing some users to enjoy a higher value from consuming a niche content than the popular content (i.e. $1 < x \leq 2$) captures these observations found in the literature.

Furthermore, we assume users do not incur any search cost and sample contents until they find one that matches their taste. This assumption is reasonable in the context of media streaming platforms and is supported empirically by the digital goods literature. Specifically, access to accurate search engines, convenient browsing and personalized recommendation systems have reduced users’ search cost and subsequently, have increased the collective share of consumption of niche and obscure products (Brynjolfsson et al., 2011).

Finally, we assume that user x’s valuations for the niche content are drawn independently from identical Bernoulli distributions with support on $\{0, x\}$ and a success probability p. One can interpret x as the maximum value a user of type x can get from consuming a niche content that matches her taste. Arguably, the IID assumption is only reasonable if the niche content are not differentiated vertically (e.g., they have the same attributes in terms of quality/genre) and also that the distribution of the valuation (x) only depends on the user type and not the content. The assumption of independent valuations is common in the bundling literature (see for example Bakos and Brynjolfsson (1999); Geng et al. (2005); Abdallah et al. (2021)). This assumption simplifies the analysis, but does not drive the results. Since we model the user’s decision for a single viewing, the user’s expected utility from subscribing to the service is solely driven by the probability that she finds at least a content to consume. By the weak law of large numbers, as the number of available niche contents on the platform (n) goes to infinity, this probability goes to 1, even when there are strong positive or negative correlations amongst the user preferences for the niche contents (Geng et al., 2005). In sections 7.1 and 7.2, we relax the IID assumption and show that our results hold in the presence of dependent and non-identically distributed valuations, respectively.

4 The Aggregator Platform

We begin our analysis by considering the case where the platform does not engage in content creation and operates strictly as a content aggregator. We introduce the first-party content creation decision later in Section 5. The game is analyzed using backward induction. First, we consider the subgame played in the second stage between the provider and the users given an arbitrary bundle of contents (which depends on the presence of the popular provider and the number of niche providers present on the platform) and establish the platform's subscription fee strategy and the users' optimal joining decision. Then, we analyze the first period subgame played between the platform and the providers and derive the platform's optimal payment structure under each contracts and the providers joining decisions. Proofs are available in the appendix.

4.1 Licensing Contract

The model includes three players: the platform, the providers and the users. The sequence of events and the players' set of decisions are as follows. At the start of the first stage, the platform selects the number of niche content providers to invite to join (n) and the offer to make to the popular provider (R). Therefore, the platform's set of decisions in the first period is (n, R). The providers observe the offer made by the platform and choose whether to join or not. The popular provider joins if the platform's offer R is weakly greater than her outside option θ. While niche providers join if they are invited given that their outside option is normalized to zero. At the start of the second stage, the platform's only decision is to announce the subscription fee (s). The users observe the bundle of content and the subscription fee (s) and choose whether to join or not. At the end of the game, demand is realized and the platform collects the subscription fees.

In the second stage a user of type x observes the subscription fee (s) and the presence of the popular and niche contents on the platform and chooses to join if her utility $U(x)$ from joining is non-negative. Consider the case where the bundle offered by the platform does not include the popular content. Then, given that there exists n niche content in the bundle, a user of type x finds at least one content that matches her taste with probability $1-(1-p)^n$ and receives a value x from consuming it. Recall that x is uniformly distributed between $[0,2]$ and that a user realizes her value for a content only upon sampling it, and therefore, her utility at the time she makes her joining decision is unknown to her. She joins if her *expected* utility from joining is non-negative. A user of type x's expected utility from subscribing is $U(x) = (1-(1-p)^n)\ \ (x) - s$. Consequently, users with $\frac{s}{1-(1-p)^n} \leq x \leq 2$ join the bundle and the demand of the service equals $D(s) = \frac{1}{2}\left(2 - \frac{s}{1-(1-p)^n}\right)$.

On the other hand, if the popular content is included in the bundle, user x's expected utility from joining the platform when $1 < x \leq 2$ is as follows:

$$U(x)_{1<x\leq 2} = (1-(1-p)^n)\ \ (x) + (1-p)^n \frac{1}{2}\ \ (1) - s \tag{3.1}$$

Observe that with probability $(1-p)^n$, user x will not find a niche content she likes and therefore, samples the popular content. She will like the popular content (receive value 1) with probability $\frac{1}{2}$.

Meanwhile, user x's expected utility from joining the platform when $0 \leq x \leq 1$

equals:

$$U(x)_{0\leq x\leq 1}=\frac{1}{2}\ \ (1)+\frac{1}{2}\ \ (1-(1-p)^n)\ \ (x)-s \tag{3.2}$$

In effect, as described in section 3.2, when the popular content is present, for any n, a user of type $1 < x \leq 2$ expects a higher utility from joining the service than does a user with $0 < x \leq 1$, and there is a discontinuity in users' utility function at $x = 1$. For a user of type $x = 1$, we break the tie in favor of the popular content. Namely, we assume a user of type $x = 1$ samples the popular content first, and if it does not match her taste, she moves on to sampling the niche content.

A user joins if her expected utility from joining is non-negative. If $s \leq \frac{1}{2}$ the user's expected utility from joining is non-negative for any x and the demand equals $D(s) = 1$; in other words, for any x we have $U(x) \geq 0$ where $U(x)$ depends on x and is given by equation 3.1 for $1 < x \leq 2$ and equation 3.2 for $0 \leq x \leq 1$. If $\frac{1}{2} < s \leq 1-(1-p)^n$, users with $x \geq \frac{s-\frac{1}{2}}{\frac{1}{2}(1-(1-p)^n)}$ (refer to equation 3.2) expect a non-negative utility from joining the platform. As such when $\frac{1}{2} < s \leq 1-(1-p)^n$, the platform's demand is given by $D(s) = \frac{1}{2}\left(2-\frac{s-\frac{1}{2}}{\frac{1}{2}(1-(1-p)^n)}\right)$. Meanwhile, when $1-(1-p)^n < s \leq 1$, users with $x \geq \frac{s-\frac{1}{2}(1-p)^n}{1-(1-p)^n}$ expect a non-negative utility from joining the platform (refer to equation 3.1) and as such demand is given by $D(s) = \frac{1}{2}\left(2-\frac{s-\frac{1}{2}(1-p)^n}{1-(1-p)^n}\right)$. Finally, when $s > \frac{1}{2}$ no users join and $D(s) = 0$.

In order to simplify the notation, hereafter, we define the variable $y = (1-p)^n$ as the probability that a user is not matched with any niche contents, given that there are n niche content available in the bundle. Note that p is a given parameter and therefore for any value of y we derive $n = \frac{\log y}{\log(1-p)}$.

The total revenue generated in the second stage is $\Pi(s) = D(s)s$ where $D(s)$ denotes the demand for the service given the subscription fee (s), which is essentially the mass of users who choose to join the platform. The platform's optimal subscription fee decision (s^*) is the argument that maximizes $\Pi(s)$ which depends on the presence of popular and niche providers on the platform, and is provided in Lemma 3.1.

Lemma 3.1. *In the pricing subgame, the optimal subscription fee depends on the the presence of providers in the following way:*

1. *If the popular provider and n niche providers are present, if $y \leq \frac{1}{2}$, the platform sets $s^* = \frac{1}{4}(3 - 2y)$, otherwise, the platform sets $s^* = \frac{1}{2}$.*
2. *If the popular provider is absent and n niche providers are present, the platform charges $s^* = 1 - y$.*

Proof. An arriving user observes the bundle of content available on the platform. If the popular content is present, a user x, with $0 \leq x \leq 1$, expects a utility $U(x) = \frac{1}{2}(1) + \frac{1}{2}\ \ (1 - y)\ \ (x) - s$ from joining the platform and a user with an $1 < x \leq 2$ expects a utility $U(x) = (1-y)\ \ (x) + y\ \ (\frac{1}{2}) - s$. For any value of $y < 1$, the expected utility of a user of type x where $0 \leq x \leq 1$ is less than that of user of type x where $1 < x \leq 2$. If $s \leq \frac{1}{2}$, then users join regardless of their taste parameter x, $D(s) = 1$ and $\Pi(s) = D(s)s = s$. The profit is strictly increasing in s and therefore, $s^* = \frac{1}{2}$ and $\Pi(s^*) = \frac{1}{2}$. If $s > \frac{1}{2}$, $D(s) = 1 - \frac{2s-1}{2(1-y)}$, and $\Pi(s) = \left(1 - \frac{2s-1}{2(1-y)}\right) s$. The profit function is concave in s. Therefore, if $y \leq \frac{1}{2}$, $s^* = \frac{1}{4}(3 - 2y)$ and $\Pi(s^*) = \frac{(3-2y)^2}{16(1-y)}$, otherwise, $s^* = \frac{1}{2}$ and $\Pi(s^*) = \frac{1}{2}$. If the popular content is not present, then a user x expects a utility equal to $U(x) = (1 - y)\ \ (x) - s$. Consequently, we have $D(s) = 1 - \frac{s}{2(1-y)}$ and $\Pi(s) = \left(1 - \frac{s}{2(1-y)}\right) s$. Since $\Pi(s)$ is concave in s, $s^* = 1 - y$ and $\Pi(s^*) = \frac{1}{2}(1 - y)$. ☐

In the first stage, the platform must choose the lump-sum payment (R) to offer to the popular provider and the subset of the niche providers to invite to join the platform (y). The platform has the option not to invite the popular provider ($R = 0$). If so, the platform's optimal strategy is to invite a large number of niche providers to join; guaranteeing that a user x finds a niche content that matches her taste (i.e. $y^* \to 0$). In this case, in the second stage, according to lemma 3.1, the the optimal subscription fee is $s^* = 1$ and subsequently, $D(s^*) = \frac{1}{2}$. The platform's expected revenue from the bundle equals $\frac{1}{2}$. Meanwhile, if the platform chooses to invite the popular provider, she must ensure R meets the popular provider's participation constraint. The platform's expected revenue is derived from the subgame perfect equilibrium that is followed in the second stage (as described in lemma 3.1. Note from 3.1 that in the presence of the popular provider and when $y \leq \frac{1}{2}$, the subgame that is followed is such that the platform sets $s^* = \frac{1}{4}(3-2y)$ and the realized demand is $D(s^*) = 1 - \frac{2s^*-1}{2(1-y)}$. Therefore, the platform's optimization problem in the first stage becomes:

$$\begin{gathered}\max_{y,R} D(s^*)(s^*) - R \\ s.t.\;\; y < \frac{1}{2} \\ R \geq \theta\end{gathered}$$

We find, the expected profit is strictly decreasing in y. Therefore, the platform sets $y^* = 0$ and offers $R^* = \theta$ to the popular provider. Replacing the parameters with their values, the expected revenue equals $\pi = \frac{9}{16} - \theta$, which brings us to the following theorem:

Theorem 3.1. *With the licensing contract, in equilibrium, the aggregator platform pays $R = \theta$ to the popular provider, if and only if the popular provider's outside option (θ) is smaller than the expected surplus generated by her presence on the platform (i.e. $\theta \leq \frac{1}{16}$). The platform invites a large pool of niche providers, regardless.*

Proof. In the first period, the platform maximizes:

$$\begin{gathered}\max_{y,r} \frac{(3-2y)^2}{16\,(1-y)} - R \\ s.t.\;\; y < \frac{1}{2} \\ R \geq \theta\end{gathered}$$

FOC gives:

$$\frac{\partial}{\partial y}\left(\frac{(3-2y)^2}{16\,(1-y)} - R\right) = -\frac{3-8y+4y^2}{16(1-y)^2} < 0$$

Therefore, $y^* = 0$, and R^* binds the popular provider's participation constraint.

As such, the platform invites the popular provider to join if and only if the revenue generated from the presence of the popular content (i.e. $\frac{9}{16} - \theta$) is greater than the revenue generated from the bundle of a large number of niche contents (i.e. $\frac{1}{2}$). $\square$

In other words, the platform invites the popular provider if and only if the platform's profit from the presence of the popular provider is greater than the platform's profit in the absence of the popular provider (i.e. $\frac{9}{16} - \theta \geq \frac{1}{2} \Rightarrow \theta \leq \frac{1}{16}$). The licensing contract serves as a benchmark, providing the maximum expected revenue that the coordinated system can achieve.

4.2 Revenue Sharing Contract

In this section we begin with a brief discussion of the case where the platform does not differentiate between providers and compensate all with the same royalty fee. We then move on to the case where the platform offers different royalty fees to heterogeneous providers. We find none of the two revenue-sharing achieve coordination and explain how price commitment strategy on users side can help the platform to coordinate the two-sided market.

Homogeneous Royalty Fee

In this model, the platform collects subscription fees from all users who choose to join the service. The platform retains a fraction $1 - r$ of the total generated revenue and splits the remainder amongst providers proportional to the consumption of their content. Therefore, the platform's set of decisions in the first period is given by (y, r). All other aspects of the game remains similar to the licensing contract described in the previous section.

Consider the final stage's subgame. Because at this stage r is already determined, the platform's revenue function is written as $\pi_m = D(s)s(1 - r)$. Note, the platform pays the same royalty fee per consumption regardless of the content that is consumed.

The platform chooses s^* similar to the strategy described in Lemma 3.1.

In the first stage, the platform has the option to set $r = 0$, in which case, the popular provider's participation constraint is not met and only niche providers. The platform's revenue equals $\frac{1}{2}$. If the platform chooses to invite the popular provider, her optimal choices of y^* and r^* maximizes her expected profit given that the popular provider's participation constraint is met. Therefore, the platform solves the following maximization problem:

$$\begin{aligned} &\max_{y,r} D(s^*)s^*(1-r) \\ &s.t. \;\; rD(s^*)s^*\kappa \geq \theta \end{aligned}$$

where, from Lemma 3.1 we have $s^* = \frac{1}{4}(3-2y)$ and subsequently $D(s^*) = 1 - \frac{2s^*-1}{2(1-y)}$. κ denotes the expected rate of consumption of the popular provider. A user of type x, where $\frac{\frac{1}{2}-y}{1-y} \leq x \leq 1$, subscribes and samples the popular content first, and if it matches her taste (with a probability of $\frac{1}{2}$), consumes it. A user of type x, where $1 < x \leq 2$, however, samples the niche contents first and if she does not find one that matches her taste (with a probability of y), she moves on to sampling the popular content. The probability that a user x who is present on the platform consumes a content is $1 - \frac{1}{2}y$. Where $1 - \frac{1}{2}y$ denotes the probability that a user finds at least one content that matches her taste amongst all content available in the bundle. Therefore, the expected consumption of the popular content is given by:

$$\begin{aligned} \kappa(s) &= \frac{\int_{\frac{1-2y}{2(1-y)}}^{1} \frac{1}{2} f(x)\,\mathrm{d}x + \int_1^2 \frac{1}{2} y f(x)\,\mathrm{d}x}{\int_{\frac{1-2y}{2(1-y)}}^{2} \left(1 - \frac{1}{2}y\right) f(x)\,\mathrm{d}x} \\ &= \frac{1+2y-2y^2}{6-7y+2y^2} \end{aligned}$$

The rate of consumption of popular content is strictly increasing in y. The platform faces a trade-off between inviting a large pool of niche content and hence increasing the net generated revenue, and reducing her share of the revenue by setting a higher r to compensate the popular provider. On the other hand, higher y entails smaller r, but also smaller total generated revenue. We find for $\theta \leq \frac{3}{304}$, the platform invites all niche providers, and sets r such that it binds the popular provider's participation constraint. However, for higher values of θ, $y^* \geq 0$ and y^* is increasing in θ.

Proof. In the first period, the platform maximizes:

$$\begin{gathered}\max_{y,r} \frac{(3-2y)^2}{16\,(1-y)}(1-r)\\ s.t.\;\; y < \frac{1}{2}\\ r\frac{(3-2y)^2}{16\,(1-y)}\frac{1+2y-2y^2}{6-7y+2y^2} \geq \theta\end{gathered}$$

Given that the platform's profit function is strictly decreasing in r, r^* is the argument that binds the popular provider's participation constraint; $r^* = \frac{\theta}{\frac{(3-2y)^2}{16(1-y)}\frac{1+2y-2y^2}{6-7y+2y^2}}$. Plugging in r^*, the platform's expected profit becomes $\frac{(3-2y)^2}{16(1-y)} - \frac{6-7y+2y^2}{1+2y-2y^2}\theta$. The optimal choice of y depends on the value of θ. Note we are interested in the θ parameter region where there exists $0 \leq y \leq \frac{1}{2}$ such that the expected revenue from inviting the popular provider is greater than that of inviting a large pool of niche providers and earning $\frac{1}{2}$. Therefore, we include the platform's participation constraint in the pursue of finding y^*. The first-order condition is:

$$\frac{\partial}{\partial y}\left(\frac{(3-2y)^2}{16\,(1-y)} - \frac{6-7y+2y^2}{1+2y-2y^2}\theta\right) = -\frac{3-8y+4y^2}{16\,(1-y)^2} + \frac{19-28y+10y^2}{(1+2y-2y^2)^2}\theta$$

Meanwhile, the second-order condition is:

$$\frac{\partial^2}{\partial y^2}\left(\frac{(3-2y)^2}{16\,(1-y)} - \frac{6-7y+2y^2}{1+2y-2y^2}\theta\right) = \frac{1}{(1-y)^3} - \frac{4\,(26-57y+42y^2-10y^3)}{(1+2y-2y^2)^3}\theta$$

For $\theta \leq \frac{1}{832}$, the expected revenue is concave upwards in y (recall $0 \leq y < \frac{1}{2}$) and has a local minimum. Therefore, the expected revenue is maximized at one of the corner solutions. Plugging in $y = 0$ and $y = \frac{1}{2}$, we find for $\theta \leq \frac{1}{832}$, the expected revenue is maximized at $y^* = 0$. Meanwhile, for $\frac{1}{832} < \theta \leq \frac{1}{16}$, the revenue function has a point of inflection at y'', below which the expected revenue is concave downwards and above which is concave upward. When $\frac{1}{832} \leq \theta \leq \frac{3}{304}$, and $y \leq y''$ the revenue function is strictly decreasing in y. While for $y > y''$, the expected revenue is strictly increasing in y. Therefore, plugging in $y = 0$ and $y = \frac{1}{2}$, we find for $\frac{1}{832} \leq \theta \leq \frac{3}{304}$, the expected revenue is maximized at $y^* = 0$. Finally, for $\frac{3}{304} \leq \theta \leq \theta^{hr}$, y^* is the solution to the first-order condition, which is $y^* < y''$, and therefore, is the local maxima of the expected revenue function for $0 < y < \frac{1}{2}$. y^* is the first root of $(3 - 2y)(1 - 2y)\left(-2y^2 + 2y + 1\right)^2 = 16\theta(1 - y)^2\left(19 - 28y + 10y^2\right)$, and θ^{hr}, is the value of θ above which it is not worthwhile for the platform to invite the popular provider($\theta^{hr} \approx 0.01$). $\square$

This result echos that of Lei and Swinney (2018), in that, a revenue-sharing contract with a single revenue-sharing parameter and a linear revenue-split rule cannot optimally allocate the revenue amongst providers. The results also highlights the fundamental inequity between popular providers' compensation and the surplus they bring to the platform, which has been a topic of legal debate in recent years. The debate is most prominent in the music streaming industry, with platforms such as Spotify and Pandora employing similar revenue-sharing structures.

Heterogeneous Royalty Fees

We now return to the model with individual revenue-sharing model, where the platform compensates the popular provider with a share of the revenue r which is specific to the popular provider. We begin by deriving the second stage's subgame perfect equilibrium. At this point in the game, the revenue sharing parameters (r) has been determined. Under revenue sharing, given that the popular provider is present, the platform pays $r\kappa$ fraction of the total generated revenue to the popular provider. The probability that a user who is present on the platform finds a content that matches her taste equals $1 - \frac{1}{2}y$. While, the probability that a user x with $0 \leq x \leq 1$ consumes the popular content equals to $\frac{1}{2}$ and a user x with $1 < x \leq 2$, consumes the popular content if and only if she searches for a content that matches her taste amongst niche contents and does not find any (with a probability of y). Thus, the expected consumption ratio of the popular content is given by:

$$\begin{aligned}\kappa(s) &= \frac{\int_{\frac{2s-1}{1-y}}^{1}\frac{1}{2}f(x)\,\mathrm{d}x+\int_{1}^{2}\frac{1}{2}yf(x)\,\mathrm{d}x}{\int_{\frac{2s-1}{1-y}}^{2}\left(1-\frac{1}{2}y\right)f(x)\,\mathrm{d}x}\\ &= \frac{\frac{1}{4}-\frac{2s-1}{4(1-y)}+\frac{1}{4}y}{\left(1-\frac{2s-1}{2(1-y)}\right)\left(1-\frac{1}{2}y\right)}\end{aligned} \tag{3.3}$$

Let π_m be the profit of the platform. The platform's optimal subscription fee is the argument that maximizes the platform's revenue given by:

$$\begin{aligned}&\max_s\left(1-\frac{2s-1}{2\left(1-y\right)}\right)s\left(1-r\frac{\frac{1}{4}-\frac{2s-1}{4(1-y)}+\frac{1}{4}y}{\left(1-\frac{2s-1}{2(1-y)}\right)\left(1-\frac{1}{2}y\right)}\right)\\ &s.t.\quad \frac{1}{2}<s\leq 1-\frac{1}{2}y\end{aligned}$$

Lemma 3.2 describes the platform's optimal subscription fee strategy, given the presence of the popular provider and r and y parameters. Lemma 3.2 entails that when the popular content is present on the platform, for a non-zero revenue-sharing parameter $r > 0$, the platform's optimal subscription fee under revenue-sharing is strictly greater than that of licensing. In effect, given that the platform must share the generated revenue with the popular provider proportional to the consumption of the popular content, she attempts to extract some of this revenue by pricing the subscription higher such that the users with unfavorable taste for the niche content do not participate.

Lemma 3.2. *There exists a unique s^* that maximizes π_m and is governed by the presence of the popular provider and niche providers on the platform:*

1. *When the popular provider is present, the optimal subscription fee depends on parameters y and r. If $0 \leq y < \frac{5-\sqrt{9-8r}}{2(2+r)}$, $s^* = \frac{(2+r)y^2-7y+2(3-r)}{4(2-r-y)}$, otherwise, $s^* = \frac{1}{2}$.*

2. If the popular is not present, $s^* = 1 - y$.

Proof. Under revenue sharing, the user's decision to join is analogous to that of Lemma 3.1. If the popular content is present, the platform's maximization problem becomes:

$$\max_s \left(1 - \frac{2s-1}{2(1-y)}\right) s \left(1 - r\frac{\frac{1}{4} - \frac{2s-1}{4(1-y)} + \frac{1}{4}y}{\left(1 - \frac{2s-1}{2(1-y)}\right)\left(1 - \frac{1}{2}y\right)}\right)$$
$$s.t. \quad s > \frac{1}{2}$$

The FOC gives:

$$\frac{\partial}{\partial s}\left(\left(1 - \frac{2s-1}{2(1-y)}\right) s \left(1 - r\frac{\frac{1}{4} - \frac{2s-1}{4(1-y)} + \frac{1}{4}y}{\left(1 - \frac{2s-1}{2(1-y)}\right)\left(1 - \frac{1}{2}y\right)}\right)\right) =$$
$$1 - \frac{s}{1-y} - \frac{2s-1}{2(1-y)} + \frac{rs}{2(1-y)\left(1-\frac{1}{2}y\right)} - r\frac{\left(\frac{1}{4} - \frac{2s-1}{4(1-y)}\right) + \frac{1}{4}y}{1 - \frac{1}{2}y}$$

While the SOC is $-\frac{2}{1-y} + \frac{r}{(1-y)\left(1-\frac{y}{2}\right)} < 0$. The profit function is concave in s. Therefore, if $0 < r < 1$ and $y \leq \frac{5-\sqrt{9-8r}}{2(r+2)}$, $s^* = \frac{(2+r)y^2 - 7y + 2(3-r)}{4(2-r-y)}$. Otherwise, $s^* = \frac{1}{2}$ and $\Pi(s^*) = \frac{1}{2}$. ☐

In the first period, the platform simultaneously chooses the mass of niche providers to invite to join (i.e. y) and the revenue-sharing parameter (i.e. r) to compensate the popular provider with. The platform chooses r^* and y^* that maximizes her profit subject to the participation constraint of the popular provider.

$$\max_{y,r} D(s^*)s^*(1 - r\kappa)$$
$$s.t. \;\; y \leq \frac{5 - \sqrt{9-8r}}{2(r+2)}$$
$$D(s^*)s^* r\kappa \geq \theta$$

We find r^* and y^* are such that for a small enough θ ($\theta \leq \frac{5}{144}$), $y^* = 0$ and the platform invites all niche providers. Moreover, for $\theta \leq \frac{5}{144}$, r^* is the argument that binds the popular provider's participation constraint. However, for larger values of θ, to meet the popular provider's participation constraint, the platform is forced to invite less niche providers and consequently, increase the expected rate of consumption of the popular content. Therefore, for larger values of θ, the number of niche contents available is decreasing in θ.

Proof. Replacing s^*, $D(s^*)$ and κ with their values from Lemma 3.2, the maximization problem becomes:

$$\max_{r,y} \frac{\left(6-7y+2y^2-r\left(2-y^2\right)\right)^2}{16(2-y)(1-y)(2-r-y)}$$

$$s.t. \quad y \leq \frac{5-\sqrt{9-8r}}{2(r+2)}$$

$$\frac{r\left(6-7y+2y^2-r\left(2-y^2\right)\right)\left(2+3y-6y^2+2y^3-r\left(2-y^2\right)\right)}{16(2-y)(1-y)(2-r-y)^2} \geq \theta$$

Before we solve for the optimal r and y, let us analyze the revenue function of the popular provider(π_{pp}), which is strictly increasing in y, and concave in r. For $y < y'$, where y' is a fixed point and is the first root of $-8+28y-27y^2+8y^3=0$, the revenue function has a maximum at r', which is the implicit solution to the first-order condition. Therefore, for a $\theta \leq \pi_{pp}(r')$, the revenue function meets θ at two points. Because the platform's expected revenue is strictly decreasing in r, we infer that r^* is the point where θ meets π_{pp} for the first time, (i.e. at r^*, $\frac{\partial \pi_{pp}}{\partial r} > 0$). While for $y > y'$, the revenue function is strictly increasing in r and the function is maximized at $r = 1$. Furthermore, we look for the optimal solution for values of θ where the expected revenue is greater than $\frac{1}{2}$; which holds if

$$r \leq \frac{4-2y-6y^2+7y^3-2y^4-\sqrt{8\left(8y-20y^2+10y^3+13y^4-17y^5+7y^6-y^7\right)}}{\left(2-y^2\right)^2}$$

It is intuitive that the maximum platform's revenue is achieved where $\pi_{pp}(r^*, y) = \theta$. Therefore, $\frac{r^*\left(6-7y+2y^2-r^*\left(2-y^2\right)\right)\left(2+3y-6y^2+2y^3-r\left(2-y^2\right)\right)}{16(2-y)(1-y)(2-r^*-y)^2} - \theta = 0$ With the envelop theorem at r^*, $\frac{d\pi_m}{dy} = \frac{\partial\pi_m}{\partial y} + \frac{\partial\pi_m}{\partial r}\frac{\partial r}{\partial y}$. We have:

$$\begin{aligned}
\frac{\partial r}{\partial y} &= -\frac{\frac{d}{dy}\left(\frac{r\left(6-7y+2y^2-r\left(2-y^2\right)\right)\left(2+3y-6y^2+2y^3-r\left(2-y^2\right)\right)}{16(2-y)(1-y)(2-r-y)^2}-\theta\right)}{\frac{d}{dr}\left(\frac{r^*(6-7y+2y^2-r^*(2-y^2))(2+3y-6y^2+2y^3-r(2-y^2))}{16(2-y)(1-y)(2-r-y)^2}-\theta\right)} \\
\frac{\partial\pi_m}{\partial y} &= \frac{1}{16}\left(-\frac{r(1-r)^3}{(2-r-y)^2}+\frac{4r}{(2-y)^2}+\frac{1-r}{(1-y)^2}-(2+r)^2\right) \\
\frac{\partial\pi_m}{\partial r} &= \frac{\left(6-7y+2y^2-r\left(2-y^2\right)\right)^2}{16(2-y)(1-y)(2-r-y)^2}-\frac{\left(2-y^2\right)\left(6-7y+2y^2-r\left(2-y^2\right)\right)}{8(2-y)(1-y)(2-r-y)}
\end{aligned}$$

We find for $\theta \le \frac{5}{144}$, $\frac{d\pi_m}{dy} < 0$ and therefore, $y^* = 0$, and subsequently, r^* is the first root of $r^3 - 4(2\theta+1)r^2 + (32\theta+3)r - 32\theta = 0$. While, for $\frac{5}{144} < \theta <\approx 0.0449$, the platform's revenue is maximized at the first order condition $\frac{d\pi_m}{dy} = 0$. $\square$

Finally, the maximum surplus generated by the presence of the popular provider on the platform is strictly smaller than $\frac{1}{16}$, which indicates that compare to licensing, under revenue-sharing the platform cannot afford higher values of θ. With sequential entry, even for small values of θ where $y^* = 0$, platform's optimal pricing strategy $s^* = \frac{3-r^*}{4-2r^*}$ in the second stage is increasing in r^*. In effect, given r^*, the platform sets a high s^* to dissuade users with unfavorable taste for niche contents from joining the platform, decreasing the rate of consumption of popular content and increasing her own share of the revenue. The popular provider correctly anticipates the platform's strategy in the second stage and demands a higher r, which in turn decreases the net generated revenue.

Heterogeneous Royalty Fees with Price Commitment

Now suppose the platform is able to commit to a subscription price when making an offer to the popular provider. Therefore, the platform makes her pricing decision in the first stage, simultaneously with her y and r decisions. If the platform chooses to invite the popular provider her maximiation problem becomes:

$$\max_{r,y,s}\left(1-\frac{2s-1}{2\left(1-y\right)}\right)s\left(1-r\frac{\frac{1}{4}-\frac{2s-1}{4(1-y)}+\frac{1}{4}y}{\left(1-\frac{2s-1}{2(1-y)}\right)\left(1-\frac{1}{2}y\right)}\right)$$
$$s.t.\quad \frac{1}{2}<s\leq 1-\frac{1}{2}y$$
$$\left(1-\frac{2s-1}{2\left(1-y\right)}\right)sr\frac{\frac{1}{4}-\frac{2s-1}{4(1-y)}+\frac{1}{4}y}{\left(1-\frac{2s-1}{2(1-y)}\right)\left(1-\frac{1}{2}y\right)}\geq\theta$$

The platform's optimal choices are $y^*=0$, $r^*=\frac{32\theta}{3}$ and $s^*=\frac{3}{4}$. The platform's expected revenue in the presence of the popular provider is $\pi_m=\frac{9}{16}-\theta$, and therefore, the platform invites the popular provider if and only if $\theta\leq\frac{1}{16}$.

Theorem 3.2. *With the revenue-sharing contract, the platform achieves coordination with heterogeneous royalty fees and price commitment.*

Proof. r^* binds the popular provider's participation constraint. Therefore, $r^*=\frac{2\theta(2-y)(1-y)}{s(2-2s-y^2)}$. Given r^*, we rewrite $\pi_m=\left(1-\frac{2s-1}{2(1-y)}\right)s-\theta$. The FOCs are given by:

$$\begin{aligned}\frac{\partial\pi_m}{\partial s}&=1+\frac{1-4s}{2-2y}=0\Rightarrow s^*=\frac{1}{4}(3-2y)\\ \frac{\partial\pi_m}{\partial y}&=\frac{s(1-2s)}{2(1-y)^2}=\frac{1}{16}\left(\frac{1}{(1-y)^2}-4\right)<0\Rightarrow y^*=0\end{aligned}$$

□

4.3 Mechanism Comparison

Our primary finding is that if price commitment is not feasible, the platform prefers the licensing contract to the revenue-sharing contract. The reason is that the licensing contract decouples the content-mix and subscription fee decisions from the compensation of the popular provider, coordinating the two-sided market. With revenue-sharing the platform's subscription fee decision and consequently, the total net revenue generated from the bundle, inevitably, is influenced by the revenue-sharing parameter r. As r increases the platform must increase the price of the bundle to compensate for the high royalty fee. Furthermore, we find that under revenue-sharing, when the outside option of the popular provider is greater than a threshold, it is no longer optimal for the platform to invite all niche providers and $y^* > 0$. Finally, we find that for small enough θs the revenue-sharing contract performs near optimal, but as θ increases, the gap between the performance of the licensing and revenue-sharing contract increases.

Next we investigate consumer surplus under each mechanism. When the market is coordinated (via licensing or revenue-sharing with price commitment), the platform sets the subscription fee $s^* = \frac{3}{4}$ and invites all niche providers to join. Therefore, the expected consumer surplus is:

$$CS = \int_{\frac{1}{2}}^{1} \left(\frac{1}{2} + \frac{1}{2}x - \frac{3}{4}\right) \frac{1}{2}\mathbf{d}x + \int_{1}^{2} \left(x - \frac{3}{4}\right) \frac{1}{2}\mathbf{d}x = \frac{13}{32}$$

Under revenue-sharing, however, the expected consumer surplus depends on θ. Given the royalty fee $r^*(\theta)$ and the mass of niche content on the platform $y^*(\theta)$, the platform's optimal strategy in the second period subgame according to Lemma 3.2 is to set the subscription fee equal to $s^*(\theta) = \frac{(2+r^*(\theta))y^{*2}(\theta)-7y^*(\theta)+2(3-r^*(\theta))}{4(2-r^*(\theta)-y^*(\theta))}$. Given s^*, the consumer surplus is given by:

$$\begin{aligned} CS &= \int_{\frac{2s^*(\theta)-1}{1-2y^*(\theta)}}^{1} \left(\frac{1}{2} + \frac{1}{2}\left(1 - y^*(\theta)\right)x - s^*(\theta)\right) \frac{1}{2}\mathbf{d}x + \\ & \int_{1}^{2} \left((1 - y^*(\theta))x + y^*(\theta)\frac{1}{2} - s^*(\theta)\right) \frac{1}{2}\mathbf{d}x \\ &= \frac{s^{*2}(1-3y^*)}{2(1-2y^*)^2} - \frac{s\left(3 - 11y^* + 8y^{*2}\right)}{2(1-2y^*)^2} + \frac{5 - 10y^{*3} + 28y^{*2} - 22y^*}{4(1-2y^*)^2} \leq \frac{13}{32} \end{aligned}$$

Therefore, the consumer surplus is maximized in a coordinated market. The next proposition determines that licensing and revenue-sharing with price commitment are both socially optimal; achieving the highest total surplus generated.

Proposition 3.1. *Licensing and revenue-sharing with price commitment are the socially optimum contract.*

5 The Hybrid Platform

Consider the setting where the platform has the opportunity to invest in producing its own content. We assume that the probability of success – defined by the content becoming popular – is an increasing function of the amount of investment devoted to its production. Specifically, consistent with the literature, we assume a quadratic cost function to capture this. When the platform selects a success probability of ϕ, it incurs a cost of $c\phi^2$, where c denotes the platform's production cost factor (Feldman et al., 2019).

In stage zero, the platform makes an investment decision ϕ. In the first stage, the platform invites the providers. The providers observe ϕ, anticipate the joining decision of other providers and the platform's pricing strategy in the final stage and choose to accept (or reject) the platform's offer. In the second and final stage, the platform's production uncertainty is resolved and the platform chooses the subscription fee to charge; observing the subscription fee and the content available in the bundle, the users choose to join (or not). To simplify the model we assume that if the production fails to become a popular content the platform does not include the output in the bundle.

The timeline is justified based on evidence from practice, where the content produced by the platform is released gradually. We also assume that the uncertainty is resolved when the platform makes its subscription fee decision. For example, a video streaming platform such as Netflix adjust its subscription fee annually upon realizing the value of the bundle to the users. Meanwhile, since in practice users can join (leave) the subscription at any point in time, we assume the users' decision to join also takes place after the platform content's type is realized.

The game is solved through backward induction starting with the subscription fee subgame.

5.1 Licensing Contract

In the final stage, one of the following subgames is realized:

1. The popular provider is present and the platform's production succeeds. In this case the platform's optimal subscription fee is $s^* = \frac{3}{4}$ and all users join $D(s^*) = 1$.

2. The popular provider is present and the platform's production fails. In this case the platform's optimal subscription fee decision follows that described in Lemma 3.1.

3. The popular provider is absent, and the platform's production succeeds. In this case, the subgame equilibrium described in Lemma 3.1 applies.

4. The popular provider is absent and the platform's production fails, in which case $s^* = 1 - y$ and $D(s^*) = \frac{1}{2}$

In the first stage, given ϕ the platform must choose the content-mix. If the platform chooses to invite the popular provider, her maximization problem becomes:

$$\max_{R,y} \phi \frac{3}{4} + (1-\phi) \frac{(3-2y)^2}{16\,(1-y)} - R$$

$$s.t. \quad R \geq \theta$$

The platform optimal strategies are $R^* = \theta$ and $y^* = 0$. Meanwhile, if the platform chooses not to invite the popular provider her maximization problem becomes:

$$\max_y \phi \frac{(3-2y)^2}{16\,(1-y)} + (1-\phi)\frac{1}{2}\,(1-y)$$

The expected profit is strictly decreasing in y and thus $y^* = 0$. Therefore, in the first stage the platform makes an offer to the popular provider if and only if $\theta \leq \frac{1}{16}(1+2\phi)$. Noteworthy is that in the presence of investment in first-party content ($\phi > 0$), the expected surplus generated by the presence of popular content increases, and the platform invites the popular provider with higher values of θ, compared to when $\phi = 0$ and the platform does not make an investment. In the first stage, given the subgame equilibrium described above, the platform's optimal ϕ maximizes:

$$\max_\phi \phi\frac{3}{4} + (1-\phi)\frac{9}{16} - c\phi^2 - \theta$$
$$s.t.\;\; \theta \leq \frac{1}{16}(1+2\phi)$$

We find if $\theta \leq \frac{1}{16} + \frac{3}{256c}$, $\phi^* = \frac{3}{32c}$ and the platform invites the popular provider to join. Otherwise, $\phi^* = \frac{1}{32c}$ and the platform does not make an offer to the popular provider, which brings us to the following theorem:

Theorem 3.3. *With the licensing contract if $\theta \leq \frac{1}{16} + \frac{3}{256c}$, in equilibrium, the platform makes an investment of $I(\phi^*)$, where $\phi^* = \frac{3}{32c}$, in the production of her first-party content, invites all niche providers and makes an offer of $R = \theta$ to the popular provider.*

Proof. Consider the case where the platform invites n niche providers and chooses to invest in its own content production. The content becomes popular with probability ϕ and the platform earns an expected revenue of $\frac{(3-2y)^2}{16(1-y)}$. Whereas with probability $1-\phi$ the production fails and the bundle generates an expected revenue of $\frac{1}{2}(1-y)$. The platform's expected profit is $\phi\frac{(3-2y)^2}{16(1-y)} + (1-\phi)\frac{1}{2}(1-y)$ and is strictly decreasing in y. Therefore, $y^* = 0$. The platform's optimal investment parameter ϕ, is the argument that maximizes:

$$\max_{\phi} \quad \left(\phi\frac{9}{16} + (1-\phi)\frac{1}{2} - c\phi^2\right)$$

Therefore, $\phi^* = \frac{1}{32c}$ and $\pi_m(\phi^*) = \frac{1}{2} + \frac{1}{1024c}$. Now consider the case where given a production decision of ϕ, the platform invites the popular provider. If the platform becomes popular then a user x, with $0 \leq x \leq 1$, expects a utility $U(x) = \frac{3}{4}(1) + \frac{1}{4} \ (1-y) \ (x) - s$ from joining the platform and a user with an $1 < x \leq 2$ expects a utility $U(x) = (1-y) \ (x) + y \ (\frac{3}{4}) - s$. Consequently, the platform's optimal subscription fee is $s^* = \frac{3}{4}$, all users join, $D(s^*) = 1$ and $\Pi(s^*) = \frac{3}{4}$. Meanwhile, if the production is unsuccessful, then the platform's optimal profit is given by $\frac{(3-2y)^2}{16(1-y)}$. Therefore, the expected profit of the platform is strictly decreasing in y:

$$\max_{y} \phi\frac{3}{4} + (1-\phi)\frac{(3-2y)^2}{16(1-y)} - \theta$$

$y^* = 0$ and the platform's maximization problem becomes:

$$\max_{\phi} \phi\frac{3}{4} + (1-\phi)\frac{9}{16} - c\phi^2 - \theta$$

Therefore, $\phi^* = \frac{3}{32c}$ and $\pi_m = \frac{9}{16} + \frac{9}{1024c} - \theta$.

Therefore, as long as $\frac{9}{16} + \frac{9}{1024c} - \theta > \frac{1}{2} + \frac{1}{1024c}$ the platform invites the popular provider and invests $I\left(\frac{3}{32c}\right)$ in the production of the popular content. $\square$

5.2 Revenue-Sharing Contract With Price Commitment

If price commitment is feasible, in the second stage, given ϕ, the platform must commit to a pricing strategy which includes the subscription fee if the production succeeds (s_p) and the subscription fee if the production fails (s_f), when inviting the providers. Therefore, in the second stage the platform chooses y, r and s_p and s_f, simultaneously. If the platform chooses to invite the popular provider, the platform maximization problem becomes:

$$\begin{aligned}
\max_{y,r,s} \phi \left(1 - \frac{4s_p - 3}{2(1-y)}\right) s_p \left(1 - r\kappa_p\right) + (1-\phi) \left(1 - \frac{2s_f - 1}{2(1-y)}\right) s_f \left(1 - r\kappa_f\right) \\
s.t. \quad \phi \left(1 - \frac{4s_p - 3}{2(1-y)}\right) s_p r\kappa_p + (1-\phi) \left(1 - \frac{2s_f - 1}{2(1-y)}\right) s_f r\kappa_f \geq \theta \\
\text{where} \quad \kappa_p = \frac{\frac{3y}{8} + \frac{3}{8}\left(1 - \frac{4s_p-3}{1-y}\right)}{2\left(1 - \frac{4s_p-3}{2(1-y)}\right)\left(1 - \frac{y}{4}\right)} \text{ and } \kappa_f = \frac{\frac{y}{4} + \frac{1}{4} - \frac{2s_f-1}{4(1-y)}}{\left(1 - \frac{2s_f-1}{2(1-y)}\right)\left(1 - \frac{y}{2}\right)}
\end{aligned}$$

If $\theta \leq \frac{1}{16}(1+2\phi)$, the platform's optimal strategy is to invite all providers to join. Thus, $y^* = 0$, $r^* = \frac{64\theta}{6+3\phi}$ and $s_p^* = s_f^* = \frac{3}{4}$. Otherwise, the platform invites all niche providers to join ($y^* = 0$), and sets $s_p^* = \frac{3}{4}$ and $s_f^* = \frac{1}{2}$. When the platform invites the popular provider, she commits to the same subscription fee regardless of the outcome of the production. However, if the production succeeds, all users join the service and $D_p = 1$. Whereas if the production fails, only users with $x \geq \frac{1}{2}$ join the platform and $D_f = \frac{3}{4}$. In other words, the addition of first-party popular content to the bundle generates revenue by increasing the user base of the platform.

Theorem 3.4. *Under the revenue-sharing contract with price commitment, if* $\theta \leq \frac{1}{16} + \frac{3}{256c}$*, in equilibrium, the platform makes an investment of* $I(\phi^*)$*, where* $\phi^* = \frac{3}{32c}$*, in the production of her first-party content, invites all niche providers and makes an offer of* $r = \frac{2048c\theta}{9+192c}$ *to the popular provider.*

Proof. Consider the case where the platform invites the popular provider and the platform itself becomes a popular provider. The platform's profit function given ϕ is given by:

$$\begin{aligned}
\max_{y,r,s} \phi \left(1 - \frac{4s_p - 3}{2(1-y)}\right) s_p \left(1 - r\kappa_p\right) + (1-\phi) \left(1 - \frac{2s_f - 1}{2(1-y)}\right) s_f \left(1 - r\kappa_f\right) \\
s.t. \quad \phi \left(1 - \frac{4s_p - 3}{2(1-y)}\right) s_p r\kappa_p + (1-\phi) \left(1 - \frac{2s_f - 1}{2(1-y)}\right) s_f r\kappa_f \geq \theta
\end{aligned}$$

Let us consider the expected consumption rate of the popular content when the first-content becomes popular. A user of type $0 \leq x \leq 1$, finds a popular content that matches her taste with probability $\frac{3}{4}$. Meanwhile, a user of type $1 < x \leq 2$ consumes one of the popular contents with a probability $y\frac{3}{4}$. Finally, a user finds at least a single content to consume with probability $1 - \frac{1}{4}y$. Since there are two popular contents available and that the user randomly chooses one of the contents to sample first, we have:

$$\begin{aligned} \kappa_p &= \frac{\int_{\frac{4s_p-3}{1-y}}^{1} \frac{3}{4} f(x)\,\mathrm{d}x + \int_{1}^{2} \frac{3}{4} y f(x)\,\mathrm{d}x}{2\int_{\frac{4s_p-3}{1-y}}^{2} \left(1 - \frac{1}{4}y\right) f(x)\,\mathrm{d}x} \\ &= \frac{\frac{3y}{8} + \frac{3}{8}\left(1 - \frac{4s_p-3}{1-y}\right)}{2\left(1 - \frac{4s_p-3}{2(1-y)}\right)\left(1 - \frac{y}{4}\right)} \end{aligned}$$

While κ_f is given by equation 3.3. Similar to licensing, $s_p^* = \frac{3}{4}$ and $s_f^* = \frac{1}{4}(3 - 2y)$. Furthermore, r^* binds the popular provider's participation constraint. Given $s_p^* = \frac{3}{4}$ and $s_f^* = \frac{1}{4}(3 - 2y)$, $\kappa_p = \frac{3(1+y)}{4(4-y)}$ and $\kappa_f = \frac{1+2y-2y^2}{6-7y+2y^2}$, the popular provider's participation constraint is as follows:

$$\begin{aligned} \phi\frac{3}{4}r\frac{3(1+y)}{4(4-y)} + (1-\phi)\frac{(3-2y)^2}{16(1-y)}r\frac{1+2y-2y^2}{6-7y+2y^2} &= \theta \Leftrightarrow \\ r^* = \frac{\theta}{\phi\frac{9(1+y)}{16(4-y)} + (1-\phi)\frac{(3-2y)^2}{16(1-y)}\frac{1+2y-2y^2}{6-7y+2y^2}} \end{aligned}$$

Replacing r^* in the platform's expected revenue function, y^* is the argument that maximizes:

$$\max_y \phi\frac{3}{4} + (1-\phi)\frac{(3-2y)^2}{16(1-y)} - \theta$$

The revenue function is strictly decreasing in y; $y^* = 0$. The rest of the proof is similar to the licensing case. $\square$

5.3 Revenue-Sharing Contract Without Price Commitment

When the popular provider is present and the platform's production succeeds $s^* = \frac{3}{4}$ and all users join $D(s^*) = 1$. Otherwise, given r, the results derived in Lemma 3.2 apply. If the popular provider is not present, the subscription fee subgame follows that of the licensing contract. In the first stage, given ϕ, the platform must choose the mix of providers to invite to join the platform. If the platform chooses to invite the popular provider, her maximization problem becomes:

$$\begin{aligned} &\max_{y,r} \phi \left(\frac{3}{4} - \frac{9r(1+y)}{16(4-y)} \right) + (1-\phi) \frac{\left((2+r)y^2 - 7y + 2(3-r)\right)^2}{16(2-y)(1-y)(2-r-y)} \\ \text{s.t.} \quad & \phi \frac{3}{4} r \kappa_p + (1-\phi)\kappa_f \frac{(3-2y-r(2-y))(2-y)}{4(1-y)(2-r-y)} \frac{\left((2+r)y^2 - 7y + 2(3-r)\right)^2}{4(2-r-y)} \geq \theta \\ & \text{where} \quad \kappa_f = \frac{2+3y-6y^2+2y^3 - r\left(2-y^2\right)}{(2-y)^2(3-2y-r(2-y))} \text{ and } \kappa_p = \frac{\frac{3y}{8} + \frac{3}{8}}{2\left(1-\frac{y}{4}\right)} \end{aligned}$$

The platform's optimal choice of r and y depends on the relationship between production success probability parameter (ϕ) and the popular provider's outside option (θ). The popular provider's expected revenue is concave in r and has a local maximum. Note a high r results in a high subscription fee in the subscription pricing subgame. Therefore, a high r is detrimental to the total revenue generated and, consequently, the provider's share of the revenue. Given ϕ, if $\theta \leq \theta_r(\phi)$, there exists a unique $r(y)$ that satisfies the popular provider's participation constraint. Given $r^*(\phi, \theta)$ and $y^*(\phi, \theta)$, the platform chooses to invite the popular provider if and only if her expected profit in the presence of the popular content is weakly greater than her expected profit in the absence of the popular content. Given the first stage subgame equilibrium strategies ($r^*(\phi, \theta)$ and $y^*(\phi, \theta)$), in stage zero, the platform chooses ϕ^* that maximizes her expected revenue. We numerically observe that compared with the coordinated market, its is optimal for the platform to invite the popular provider for smaller ranges of c and θ. For larger values of θ, the platform invests more in the production of first-party content. Moreover, for higher values of θ, the platform invites significantly fewer niche providers to increase the rate of consumption of popular content and mitigate the negative effects of the high royalty fee.

5.4 Mechanism Comparison

The comparison between the revenue allocation rules highlights three main results. 1) Similar to the base model, if commitment is not feasible, the licensing contract coordinates the market, and if commitment is feasible, both revenue-sharing and licensing contracts can coordinate the market. 2) The platform's investment in producing first-party content increases the expected revenue of the market, regardless of the presence of the popular provider. The amount of investment, however, depends on the price of the popular content. The platform invests significantly more in producing first-party content, if it allows her to attract the popular provider. 3) If price commitment is not feasible, with revenue-sharing the platform invests more in the production of first-party content but is worse off in terms of the expected revenue generated.

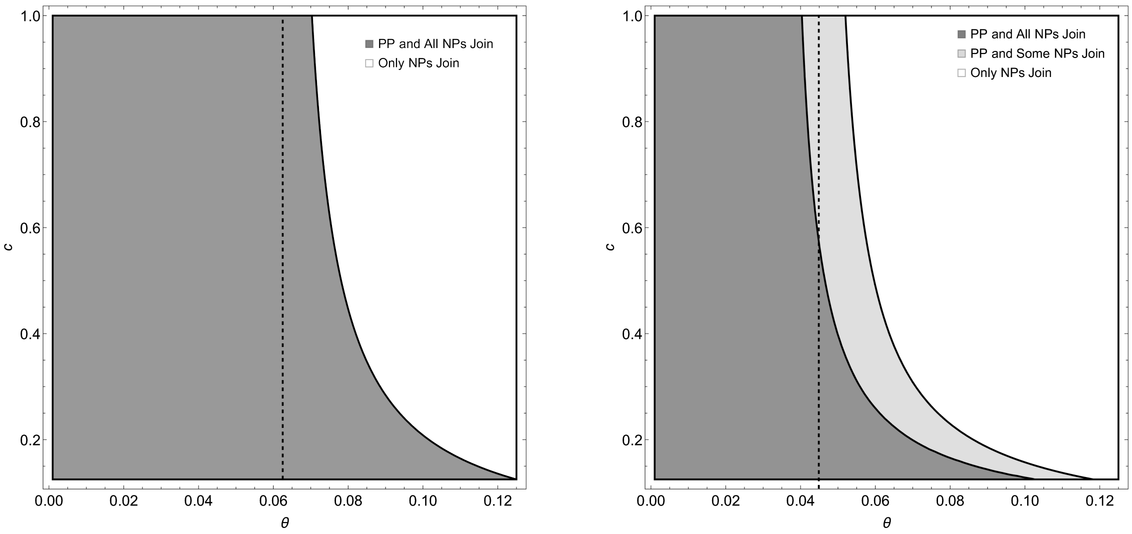

(a) Coordinated Market

(b) Revenue-Sharing without Price Commitment

Figure 3·2: Platform's optimal content-mix strategy as a function of popular provider's outside option θ and the platform's production cost c.

Figure 3·2 graphs the platform's optimal content-mix strategy as a function of model parameters c and θ with and without price commitment. In a coordinated market (Figure 3·2a), the platform's optimal content-mix includes a large pool of niche contents, but the presence of the popular content, depends on the production cost of the platform and the outside option of the popular provider, such that the platform invites the popular provider if and only if $\theta \leq \frac{1}{16} + \frac{1}{128c}$. In effect, compared to the base model, in the presence of first-party content production, the platform is able to afford higher prices demanded by the popular provider. This is because users with an unfavorable taste for niche contents value the bundle of the third-party and first-party popular content more than a single popular content. Therefore, by inviting the popular provider and investing in its own content the platform attract more users and generates more revenue. With revenue-sharing, when price commitment is not feasible (Figure 3·2b), for mid-range values of c and θ, the platform's optimal content-mix policy is to include fewer niche contents in the bundle (i.e. $y^* > 0$) to manipulate the rate of consumption of popular content. Note with $y^* > 0$ users with a favorable taste for niche contents may not find a niche content to consume with probability y^* and move on to sample the popular content. Therefore, as y^* increases the expected consumption of the popular content increases. Given higher consumption rates the platform can afford to invite the popular provider with smaller royalty fees. As shown in Figure 3·3, as θ and c increase, y^* increases, indicating that the higher is the cost of popular content acquisition, the fewer niche contents are present on the platform. However, as evident in Figure 3·2b and Figure 3·3, for small values of θ and c the platform invites all niche providers to join. Note, the smaller the y^*, the more niche providers are present on the platform.

To measure the relative performance of the revenue-sharing contract without price commitment compared to a coordinating contract from the perspective of the platform, we compute the ratio of the optimal expected profit under the two mechanisms. Thus, the smaller is the ratio, the more inefficient the revenue-sharing contract with-

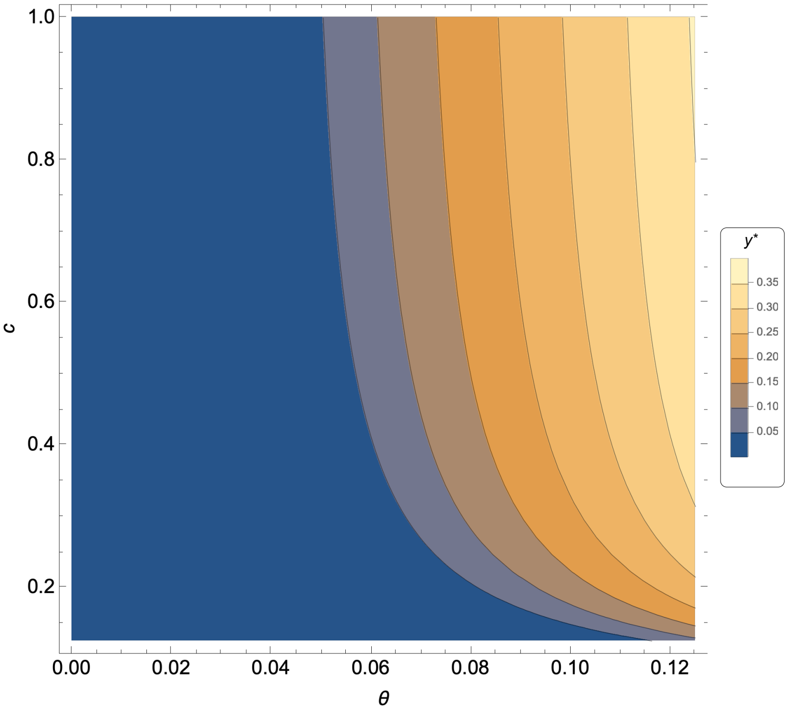

Figure 3·3: Platform's optimal content-mix strategy as a function of popular provider's outside option θ and the platform's production cost c.

out price commitment becomes. As shown in Figure 3·4, the contract inefficiency is increasing in the outside option θ. However, smaller production cost of the platform mitigates the negative effect of popular provider's higher outside option. This can be explained as follows. With a smaller production cost, the platform can invest more to ensure the success of its production and increase the expected surplus generated from the bundling of her popular content with the third-party popular content. The greater investment allows the platform to 1) satisfy the popular provider's participation constraint with a smaller royalty fee, and 2) to invite all niche providers to join, both of which ensures that the pricing and market size remain comparable to the coordinated market. As c increases the performance of the contract deteriorates rapidly.

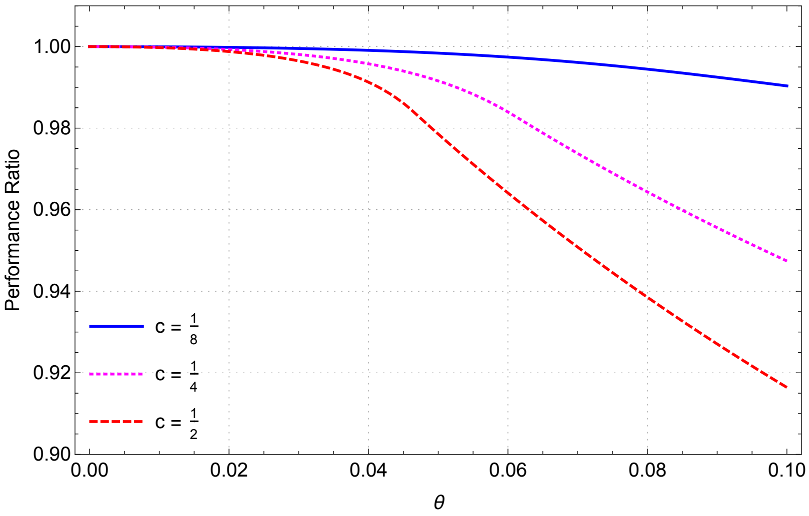

Figure 3·4: Ratio of Expected profit of the revenue - sharing contract without price commitment to expected profit of the coordinated market, as a function of popular provider' s outside option θ at different values of production cost c.

The platform's optimal investment decision ϕ can be seen in Figure3·5. For small values of θ ($\theta = 0.01$ in Figure 3·5), the platform's optimal investment decision is analogous to that of a coordinated market. However, as θ increases the platform invests more in producing its own content to be able to afford the provider's outside option.

6 Extensions

6.1 Positive Outside Options of Niche Providers

In this section, we extend the analysis by allowing niche content providers to have a positive outside option. Suppose there exists a large pool of niche providers. The outside option is homogeneous amongst niche providers and is denoted by θ_l.

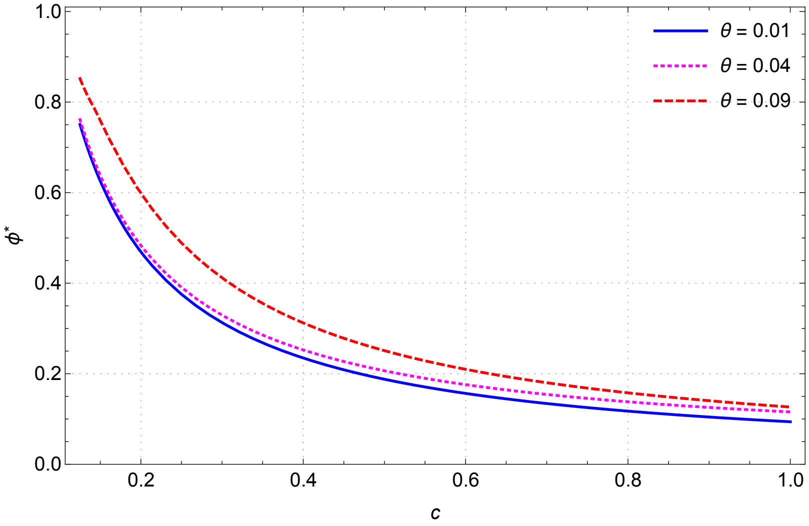

Figure 3·5: Platform's optimal investment parameter ϕ with the revenue-sharing contract without price commitment as a function of production cost c, at different values of θ

The Aggregator Platform

In the absence of the popular provider, in the second stage, given the presence of n niche providers, the platform's optimal subscription fee equals $s^* = 1 - y$ and $D(s^*) = \frac{1}{2}$. In the first stage, the platform's optimal choice of y is the argument that maximizes her expected revenue, given that the niche provider's participation constraint is met. Under licensing, the platform's maximization problem is given by:

$$\max_{n} \left(\frac{1}{2}(1-y) - \frac{\log(y)}{\log(1-p)}\theta_l \right)$$

The platform's profit function is concave in y and is maximized at $y^* = 2a$, where $a = -\frac{\theta_l}{\log(1-p)}$.

In the presence of the popular content provider, the positive outside option does not affect the pricing strategy of the platform in the final stage, as described in Lemma 3.1 and Lemma 3.2. Under licensing, in the second stage, given the presence of the popular provider, the platform's optimal strategy y^* is the argument that maximizes her expected profit given by:

$$\max_{y} \frac{(3-2y)^2}{16(1-y)} - \frac{\log(y)}{\log(1-p)}\theta_l - \theta$$
$$s.t. \quad 0 \leq y \leq \frac{1}{2}$$

For small enough values of θ_l, the platform's expected revenue has a local maxima at y^*, where y^* is the first root of polynomial $4y^3 - 8(1+2a)y^2 + (3+32a)y - 16a = 0$. Finally, given y^*, the platform makes an offer θ to the popular provider to join if and only if $\frac{(3-2y^*)^2}{16(1-y^*)} + a\log(y^*) \geq \frac{1}{2}(1-2a) + a\log(2a)$.

Under revenue sharing, in the final stage, given revenue sharing parameters r and r_l, the platform's optimal subscription fee is the argument that maximizes:

$$\max_{s} \left(1 - \frac{2s-1}{2(1-y)}\right) s\left(1 - r\kappa - r_l\kappa_l\right)$$
$$s.t. \quad \frac{1}{2} \leq s \leq \frac{1}{2}(3-2y)$$
$$\kappa = \frac{\left(\frac{1}{4} - \frac{2s-1}{4(1-y)}\right) + \frac{y}{4}}{\left(1 - \frac{2s-1}{2(1-y)}\right)\left(1 - \frac{y}{2}\right)}$$
$$\kappa_l = \frac{\left((1-y)\left(\frac{1}{4} - \frac{2s-1}{4(1-y)}\right) + \frac{1-y}{2}\right)}{\left(1 - \frac{2s-1}{2(1-y)}\right)\left(1 - \frac{y}{2}\right)}$$

κ and κ_l denote the share of the consumption of the popular content and niche contents on the platform, respectively. The platform's revenue function is concave in s and has a local maxima at $s^* = \frac{6-7y+2y^2-r\left(2-y^2\right)-r_l\left(4-7y+3y^2\right)}{4(2-r-y-r_l(1-y))}$. In the first stage, the platform's maximization problem is given by:

$$
\begin{aligned}
\max_{y,r_l,r} \ & \left(1-\frac{2s^*-1}{2(1-y)}\right) s^* \left(1 - r\kappa(s^*) - r_l\kappa_l(s^*)\right) \\
s.t. \quad & \left(1-\frac{2s^*-1}{2(1-y)}\right) s^* r_l \kappa_l(s^*) \geq -a\log(y) \\
& \left(1-\frac{2s^*-1}{2(1-y)}\right) s^* r\kappa(s^*) \geq \theta
\end{aligned}
$$

It is intuitive that licensing coordinates the market while revenue-sharing without price commitment does not. Furthermore, y^* is strictly increasing in θ_l under both mechanisms, albeit with a larger marginal rate under revenue-sharing. As in the case with the main analysis, the performance gap between the two mechanisms is increasing in θ. Interestingly, our numerical experiments show that for a given θ, the performance gap between licensing and revenue-sharing is strictly decreasing in a (Figure 3·6).

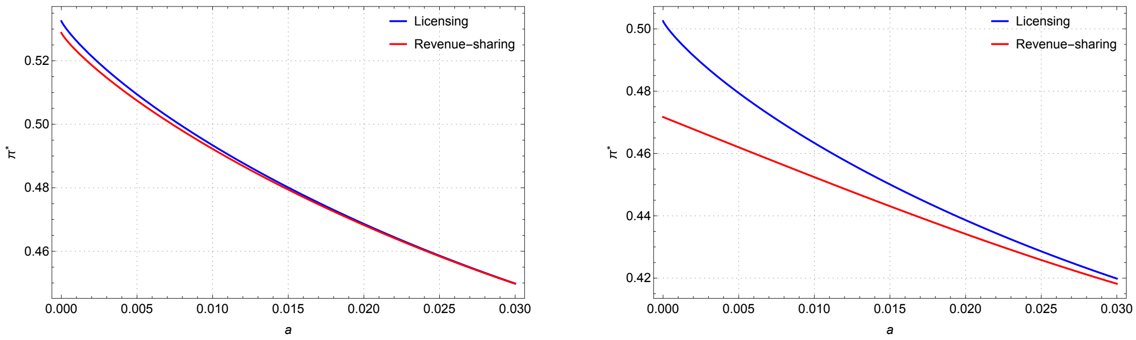

(a) $\theta = 0.03$ **(b)** $\theta = 0.06$

Figure 3·6: Platform's optimal profit as a function of niche providers outside option a, at different values of popular provider's outside option θ.

Hybrid Platform

We begin by analyzing the equilibrium in the absence of the popular provider. Upon the realization of the outcome of platform's production, if the first-party content is popular, the platform sets $s^* = \frac{1}{4}(3-2y)$, otherwise, $s^* = (1-y)$. In the second stage, the platform's optimal niche content strategy y^* is the argument that maximizes her profit, given by:

$$\max_{y} \phi \frac{(3-2y)^2}{16(1-y)} + (1-\phi)\frac{1}{2}(1-y) + a\log(y)$$
$$s.t. \quad 0 \leq y \leq \frac{1}{2}$$

The unique optimal y^* depends on a and ϕ, and is the first root of $16a - y(16 + 32a - 13\phi) + 8y^2(4+2a-3\phi) - 4(4-3\phi)y^3 = 0$. Given $y^*(\phi)$, in the first period, the platform's optimal investment strategy ϕ^* maximizes:

$$\max_{\phi} \phi \frac{(3-2y^*(\phi))^2}{16(1-y^*(\phi))} + (1-\phi)\frac{1}{2}(1-y^*(\phi)) + a\log(y^*(\phi)) - c\phi^2$$

The platform makes an offer to the popular provider if and only if θ is weakly smaller than the surplus generated by her presence.

Under revenue sharing, when the popular provider is present and the first-party content is realized to be popular, the platform's optimal subscription fee is $s^* = \frac{3}{4}$. Given s^*, $D(s^*) = 1$, and all users join. The probability that a user finds a content to consume equals $1 - \frac{1}{4}y$. The rate of consumption of the third-party popular content equals $\frac{\frac{1}{2}\left(\frac{3}{4}+\frac{3}{4}y\right)}{2\left(1-\frac{y}{4}\right)}$ and the rate of consumption of niche contents equals $\frac{\frac{1}{2}\left(\frac{1}{4}(1-y)+1-y\right)}{\left(1-\frac{y}{4}\right)}$. Meanwhile, if the production fails, the platform's optimal subscription fee is described in the previous section. ϕ, y, r and r_l jointly maximize the platform's expected revenue given by:

$$\max_{\phi,y,r_l,r} \phi \frac{3}{4}\left(1 - r\frac{3\,(1+y)}{4\,(4-y)} - r_l\left(\frac{5(1-y)}{2\,(4-y)}\right)\right)$$
$$+(1-\phi)\left(1 - \frac{2s^*-1}{2(1-y)}\right)s^*\left(1 - r\kappa - r_l\kappa_l\right)$$

where,

$$\kappa = \frac{\left(\frac{1}{4} - \frac{2s^*-1}{4(1-y)}\right) + \frac{y}{4}}{\left(1 - \frac{2s^*-1}{2(1-y)}\right)\left(1 - \frac{y}{2}\right)},$$

$$\kappa_l = \frac{\left((1-y)\left(\frac{1}{4} - \frac{2s^*-1}{4(1-y)}\right) + \frac{1-y}{2}\right)}{\left(1 - \frac{2s^*-1}{2(1-y)}\right)\left(1 - \frac{y}{2}\right)}$$

and

$$s^* = \frac{6 - 7y + 2y^2 - r\left(2 - y^2\right) - r_l\left(4 - 7y + 3y^2\right)}{4\left(2 - r - y - r_l(1-y)\right)}.$$

Figure 3·7 displays the numerical comparisons of the platform's expected earning under licensing and revenue-sharing in the presence of first-party content production. Similar to the base model, when the outside option of the niche providers is positive, licensing coordinates the market. Under licensing, y^* is independent of the outside option of the popular provider, and is strictly increasing in the outside option of niche providers. However, under revenue-sharing without price commitment, y^* is increasing in both niche providers' outside option θ_l and the popular provider's outside option θ (Figure 3·8a).

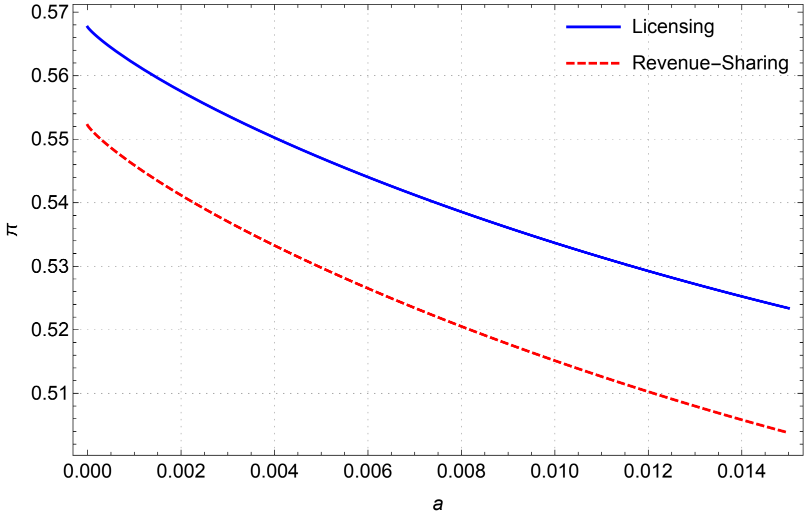

Figure 3·7: Platform's optimal expected revenue as a function of niche providers' outside option a. Parameter values: $\theta = 0.03$ and $c = \frac{1}{4}$.

Finally, similar to the main analysis, with licensing, for values of θ where the optimal choice of the platform is to invite the popular provider to join, the optimal investment is independent of the outside options of the providers. However, with revenue-sharing, the platform's optimal investment is nonmonotone in a; namely, it is initially decreasing and then increasing (see Figure 3·8b).

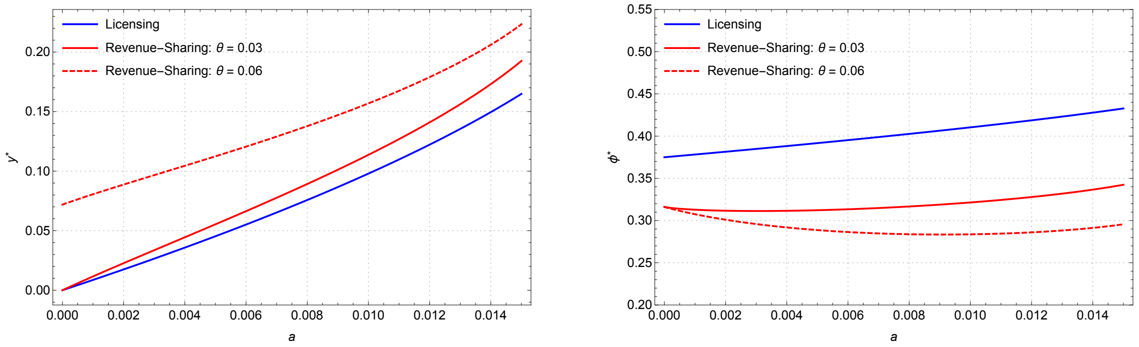

(a) Optimal y **(b)** Optimal ϕ

Figure 3·8: Platform's optimal mass of niche contents y^* and investment in first-party content parameter ϕ as a function of niche providers outside option a.Parameter value $c = \frac{1}{4}$.

7 Concluding Remarks

In this paper we study the coordination problem of video-streaming platforms under two prevalent revenue allocation practices: revenue-sharing and licensing. Video-streaming platforms such as Netflix aggregate contents from many providers and offer the bundle for a fixed fee. The platform collects the revenue generated from the subscription fees and allocates a portion of the revenue amongst providers. The platform may choose to distribute the revenue amongst providers proportional to the rate of consumption of their content (revenue-sharing) or a fixed lump-sum (licensing). We show when the providers are heterogeneous in their quality and popularity amongst users, revenue-sharing does not coordinate the market. The reason is that the platform is inclined to manipulate the consumption rate of the popular content to decrease the payment to the popular providers. However, we show that if the platform commits to a subscription fee, the revenue-sharing contract can coordinate the market. Furthermore, we find licensing is able to coordinate the market, because the platform pre-commits to a payment to the providers before setting the subscription fee. Finally, we study the platform's investment in the risky production of first-party content. We show investing in first-party content allows the platform to attract providers and users. We show, the platform invests more aggressively, to be able to attract the popular provider to join the platform.

References

Abdallah, T., Asadpour, A., and Reed, J. (2021). Large-scale bundle-size pricing: A theoretical analysis. *Operations Research*.

Alaei, S., Makhdoumi, A., and Malekian, A. (2019). Optimal subscription planning for digital goods. *Available at SSRN 3476296*.

Anderson, C. (2006). *The long tail: Why the future of business is selling less of more*. Hachette Books.

Anderson Jr, E. G., Parker, G. G., and Tan, B. (2014). Platform performance investment in the presence of network externalities. *Information Systems Research*, 25(1):152–172.

Armstrong, M. and Weeds, H. (2007). Programme quality in subscription and advertising-funded television. *Unpublished working paper*.

Armstrong, M. and Wright, J. (2007). Two-sided markets, competitive bottlenecks and exclusive contracts. *Economic Theory*, 32(2):353–380.

Bakos, Y. and Brynjolfsson, E. (1999). Bundling information goods: Pricing, profits, and efficiency. *Management science*, 45(12):1613–1630.

Bhargava, H. K. (2021). Bundling for flexibility and variety: An economic model for multiproducer value aggregation. *Management Science*, 67(4):2365–2380.

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

Caillaud, B. and Jullien, B. (2003). Chicken & egg: Competition among intermediation service providers. *RAND journal of Economics*, pages 309–328.

DeValve, L. and Pekeč, S. (2016). Optimal price/advertising menus for two-sided media platforms. *Available at SSRN 2715531*.

Feldman, P., Papanastasiou, Y., and Segev, E. (2019). Social learning and the design of new experience goods. *Management Science*, 65(4):1502–1519.

Geng, X., Stinchcombe, M. B., and Whinston, A. B. (2005). Bundling information goods of decreasing value. *Management science*, 51(4):662–667.

Guardian, T. (2018). Taylor swift strikes a blow for fellow artists as digital revenues soar. *The Guardian*.

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

Hagiu, A. and Spulber, D. (2013). First-party content and coordination in two-sided markets. *Management Science*, 59(4):933–949.

Hagiu, A. and Wright, J. (2015). Marketplace or reseller? *Management Science*, 61(1):184–203.

Jiang, H., Shilin, Z., and Zhengang, S. (2017). Two-sided market pricing in operations management: Review of current literature and research directions. *Management Science and Engineering*, 11(4):30–35.

Lei, Z. and Swinney, R. (2018). Revenue sharing and subscription platforms for digital content distribution. *Available at SSRN 3314345*.

Li, S., Luo, Q., Qiu, L., and Bandyopadhyay, S. (2020). Optimal pricing model of digital music: Subscription, ownership or mixed? *Production and Operations Management*, 29(3):688–704.

Parker, G. G. and Van Alstyne, M. W. (2005). Two-sided network effects: A theory of information product design. *Management science*, 51(10):1494–1504.

Rochet, J.-C. and Tirole, J. (2006). Two-sided markets: a progress report. *The RAND journal of economics*, 37(3):645–667.

Sanchez-Cartas, J. M. and Leon, G. (2019). Multisided platforms and markets: A literature review. *SSRN Electron. J*, pages 1–62.

Tan, T. F., Netessine, S., and Hitt, L. (2017). Is tom cruise threatened? an empirical study of the impact of product variety on demand concentration. *Information Systems Research*, 28(3):643–660.

Times, T. N. Y. (2018). Netflix will keep ‘friends’ through next year in a $100 million agreement. *The New York Times*.

Weyl, E. G. (2010). A price theory of multi-sided platforms. *American Economic Review*, 100(4):1642–72.

Appendix

7.1 Correlated Tastes for Niche Contents

Consider a bundle of niche contents $1, 2, \ldots j, \ldots, N$, witch zero marginal costs, available on the platform. When a user arrives, she randomly chooses a content to sample. If the content matches her taste, she consumes it and leave, otherwise, she moves on to the next content and repeats the process until she has sampled all contents. The user consumes content j if it matches her taste while all contents tested earlier do not. The user does not match with a content with an unconditional probability of $q = 1-p$ and each pair of contents have a pairwise failure correlation of ρ. Intuitively, as the user samples more and more content without being matched, the conditional probability with which she will be matched with the next content she samples decreases. Furthermore, we assume the failure correlation between content $j+1$ and content $j+2$ remains equal to ρ regardless of the number of failed tries among the j contents the user has already sampled. Define z_j to be a random variable, representing the user's taste for content j after she has sampled it, where $z_j = 1$, denotes that the content has not matched with the user's taste. Suppose the user has not been matched with the last $n-1$ contents she has sampled, then the conditional probability that she will not be matched with content n is written as:

$$q_n = E\left(z_n | z_1 = 1, z_2 = 1, \ldots, z_n - 1 = 1\right)$$

The pairwise correlation ρ and the unconditional failure probability q, imply $q_1 = q$, $q_2 = q + (1-q)\rho$ and q_n can be computed sequentially and equals $q_n = 1 - (1 - q)(1-\rho)^{n-1}$. Therefore, the conditional probability of N failures (the probability that the user cannot find a content, that matches her taste, to consume) in a bundle of N niche content equals:

$$P[N] = \mathbb{E}\left[\prod_{n=1}^{N} q_n\right]$$

Figure 3·9 simulates $P[N]$ as a function of N and ρ. As shown, when $\rho = 0$, $P[N] = q^N$, which is analogous to the base model. As, ρ increases, however, the probability that the user is not matched with any content increases. In other words, with greater values of ρ, the more the user samples contents and does not find a content that matches her taste, the less likely it is that she will not find such a content with more sampling. Therefore, it is intuitive that the greater the ρ, the smaller the user's expected utility from joining the platform and the smaller the subscription fee that the platform can charge. Also, the marginal benefit of the presence of the popular content increases with ρ.

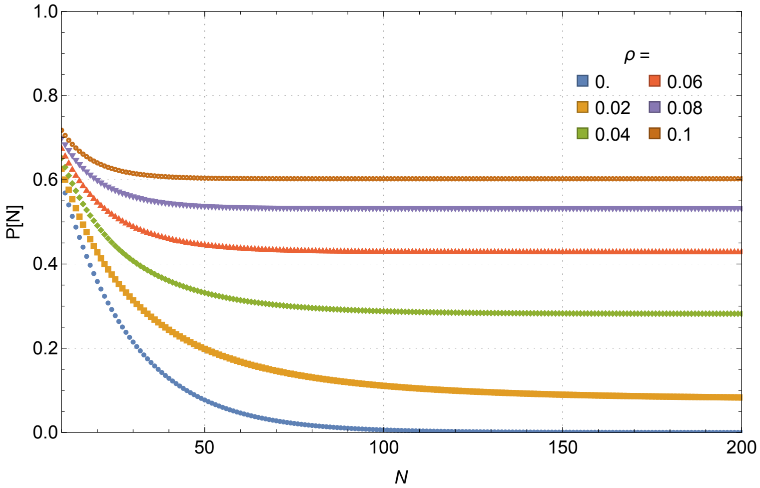

Figure 3·9: $P[N]$ as a function of N and ρ, $q = 0.95$

7.2 Non-Identical and Independent Tastes for Niche Contents

Suppose, the user's probability of being matched for contents available in the bundle is a sequence of independent Bernoulli trials and assume that the probability that the i^{th} content is matched with the user's taste is exponentially decreasing in i, with a factor λ. The user is aware of her preference ranking and λ and searches the bundle for a content that matches her taste, optimally. Given a bundle of size N, the probability that the user finds a content that matches her taste equals:

$$P[\text{success} > 0] = 1 - \prod_{n=1}^{N} \left(1 - \exp^{-\lambda n} p\right)$$

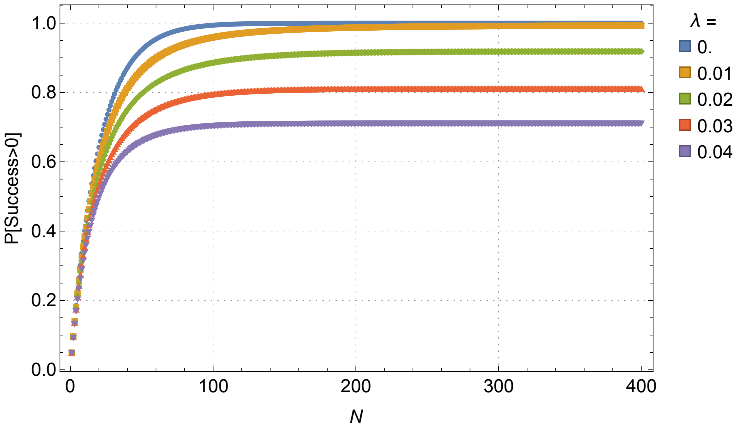

Figure 3·10: Probability of at least one success as a function of N and λ, $p = 0.05$

Figure 3·10 compares the probability of at least one success given different values for λ. It is intuitive that as λ gets larger, then the user is likely to be matched with fewer contents available on the platform. In the case of rapidly declining preferences, the platform is strictly worse off with a bundling business model, compared to selling the contents individually. However, when λ is small and the probability of success of subsequent contents decreases but slowly, then a large bundle of content can almost surely guarantee that a user finds a content to consume. Therefore, for small values of λ the optimality of bundling, and subsequently the results derived in previous sections hold.

Chapter 4

First-Party Bias in Media Streaming Platforms

1 Introduction

An important function of most video streaming platforms is the integration of sophisticated search engines and recommendation systems that provide users with personalized content recommendations. From the users' perspective, personalized recommendation systems make the content search process convenient and less time-consuming. Often, streaming services carry a large number of content in their libraries, and recommendation services offer an opportunity to expose users to content that best matches their tastes amongst many alternatives. Therefore, a large user segment of such platforms relies on the recommendation of engines to choose what content to consume. For example, according to reports published by Netflix, 75% – 80% of viewer activity on its platform is influenced by its recommendation algorithm[1].

However, the goal of a platform's recommendation algorithm is not necessarily to maximize its users' utility through recommending the most relevant content. Instead, the platform may be inclined to steer users toward consuming a content with the objective of increasing the overall profitability of the platform. The incentive of the platform to manipulate its recommendation algorithm is even higher when the platform offers first-party content. Specifically, the platform may opt to bias its advice

---

[1] How Netflix Uses Analytics To Select Movies, Create Content, and Make Multimillion Dollar Decisions [Internet]. Available from: https://blog.kissmetrics.com/how-netflix-uses-analytics/.

in favor of its own content, rather than those of third-party providers. For example, Netflix's algorithm recommended House of Cards, its first production hit, to most of its users, regardless of their personal preferences and past behavior[2]. Own-content recommendation bias is not limited to media streaming platforms. Google and Amazon, for instance, were investigated for using their search engines to promote their own products (De Corniere and Taylor, 2019). The downside to promoting own content is lower service quality. Specifically, Bourreau and Gaudin (2018) show that users, aware of the platform's recommendation bias, consider the accuracy of recommendations; and as a result, their willingness to pay for the service will be lower. Therefore, platforms face a trade-off between maximizing their profit from the consumption of more profitable content, and increasing demand by ensuring a higher quality of service. The purpose of this paper is to explore the design of personalized recommendation services and the interaction with the users' search amongst alternative options. Our specific research questions are as follows.

- How does the design of recommendation services affects users' optimal search behavior?
- How do the users' search costs affect the platform's optimal recommendation strategy?
- How does the exogenous third-party provider's royalty fee affects the platform's recommendation strategy?

Consistent with the body of work on information intermediaries and search engine bias in economics (Hagiu and Jullien, 2011; Armstrong and Zhou, 2011), we define strategic recommendation bias as the deliberate decision of the platform to steer a

---

[2] https://www.nytimes.com/2013/02/25/business/media/for-house-of-cards-using-big-data-to-guarantee-its-popularity.html

user away from their ideal content. We develop a game-theoretical model featuring one platform, one third-party content provider, and users. The platform offers two vertically differentiated contents, one supplied by the platform and the other by the third-party provider, in a bundle to users for a subscription fee. Users are heterogeneous in their preferences for the two contents. We use a similar design for consumer search as that of Armstrong et al. (2009). Namely, we consider users *a priori* have imperfect information about their preferences for each content available on the platform and realize their valuation by sequentially sampling the content at a cost. Furthermore, we consider that users sample the first content at no cost, but incur a search cost if they choose to sample the second content. We examine a setting where the platform has perfect information about a user's valuation for each product and provides a personalized recommendation to each user. In practice, sophisticated recommendation systems learn about a user's preferences based on 1) the user's viewing history and how the user rated other contents, 2) viewing history of other users with similar tastes and preferences and 3) information about the contents, such as genre, categories, actors, release year, etc[3]. This allows the platform to gain precise information about users' preferences for new content. The platform may manipulate the search process by steering users to the first-party content. Our model sheds light on how revenue-sharing structures and users' strategic response to the personalized recommendations they receive, determine platforms' strategic design of personalized recommendation services provided to their users.

A strategic user observes the platform's recommendation and samples the recommended content at no cost. She realizes her valuation for the content, and updates her belief about her valuation for the alternative content, accordingly. Given her updated belief, the user may choose to sample the alternative content for a cost. Therefore, the platform's optimal recommendation strategy incorporates the strate-

---

[3] https://help.netflix.com/How Netflix's Recommendations System Works

gic user's searching behavior. To the best of our knowledge, this is the first paper that incorporates users' learning and search behavior in designing a platform's optimal recommendation strategy. We find that the dynamics of users search behavior in the presence of recommendation bias, impacts the platform's optimal strategy such that 1) the platform' strategic bias in favor of first-party content is weakly increasing in the users' search cost and the third-party provider's royalty rate; 2) the platform's revenue is weakly decreasing in the third-party provider's royalty rate, but weakly increasing in the user's search cost; 3) social welfare is weakly decreasing in the platform's recommendation bias.

2 Literature Review

This paper relates to literature on personalized recommendation systems. There is a large body of literature that focuses on the technical aspects of developing personalized recommendation systems that maximize users' probability of purchase (see Gaur and Liu (2020) for a recent overview). A few papers have focused on developing recommendation algorithms with the objective to maximize the platform's profit instead of user's purchasing probability (Choi and Mela, 2019; Choudhary and Zhang, 2019; Dinerstein et al., 2018). Most recently, Zhou and Zou (2021) study a marketplace that recommends products which lead to the highest expected profit. They show that in such settings, third-party sellers are incentivized to adjust their prices to compete for recommendations, which, ultimately results in decreasing the platform's profit. However, there is limited understanding of how a user's search behavior is affected when the recommendation is designed to maximize the platform's profit and not the user's utility. As such, our paper extends this literature by investigating the effects of recommendation systems on users' search behavior and market outcome.

Our work relates to the literature on information intermediaries. In many settings, similar to streaming platforms, service providers have access to more information than their customers. As such, prior work has investigated the intermediary's choice to disclose information or recommendations that affect buyers' behavior. Hagiu and Jullien (2011) study an online intermediary's search diversion strategy. They show that an intermediary has an incentive to lower the quality of their recommendation in exchange for higher revenues. They also identify conditions under which search diversion is not beneficial to the intermediary. However, they assume that, if users choose to participate (i.e., subscribe), they always follow the recommended action of the intermediary. In this paper, however, we consider users to be strategic in their response to content recommendations, and study how the design of the recommendation systems is influenced by users' search behavior. Papanastasiou et al. (2018) study the intermediary's optimal information provision design and show that partial information structures can maximize aggregate consumer surplus by influencing the purchasing decisions of users. Other papers in the economic literature have looked at "intermediation bias", as the intermediary platform's choice to utilize its technology to "direct" user behavior. De Corniere and Taylor (2019) study the determinants of intermediation bias and its consequences. They consider a vertically-integrated intermediary that biases its recommendation in favor of its subsidiary seller at the expense of third-party sellers. Hagiu et al. (2020) studies the case of Amazon and show how the platform may engage in preferencing its own product through its recommendation system which effectively damages third-party sellers and raises antitrust concerns.

This work also complements the broader strand of literature that studies the the design and optimization of two-sided digital goods platforms. Hagiu and Wright (2015) study the trade-offs faced by a two-sided intermediary that facilitates transactions between buyers and sellers and must choose to operate as a marketplace or a reseller. Other papers explore the implications of specific contract forms in the selling of digital content (e.g. the agency model and wholesale contracts (Johnson, 2017)) and which channel to buy or sell from (e.g. content providers decision to sell a la carte or via subscription (Lei and Swinney, 2018)). Feldman et al. (2019) study whether food-delivery platforms are beneficial to restaurants. Bimpikis et al. (2020) examine the mechanisms through which information design on supply-side decisions increases platform revenues. The current paper considers a two-sided platform setting, but the focus is on the implications of the design of personalized recommendation systems for the user side of the market. In investigating these implications, we model users as strategic agents who make their own content choice, and we show that the design of personalized recommendations can help the platform increase its revenues by influencing the users' search behavior. Finally, our paper also broadly relates to the literature on the bundle and subscription business model. Cachon and Feldman (2011) demonstrate that users become less price sensitive when considering a series of consumption opportunities rather than considering them individually. Therefore, subscription pricing is capable of extracting more revenue from customers than peruse pricing. Lei and Swinney (2018) show inducing high quality content provider to join a subscription platform is challenging and costly, because a revenue-sharing contract cannot compensate high quality creators for their contribution to consumer utility.

3 Model Description

A monopolist platform offers two contents in a bundle, one supplied by herself (denoted f) and one supplied by a third-party provider (denoted t), for a fixed fee s. There are a continuum of users with a total mass normalized to one. Each user wishes to consume one content from the bundle. Nature draws each user's value for each content, randomly, from a uniform distribution on the interval $[0, 1]$. Prior to subscribing to the bundle, users have imperfect information about their valuation for each content and sequentially sample the contents to learn their value. They pay a search cost c to sample a content and can only consume a content they have sampled. Moreover, we assume users costlessly return to consume a content that best matches their tastes after they have sampled both contents. This assumption is generally imposed in the consumer search literature (Armstrong et al., 2009). Note that, ex ante, users do not have any preference for either content and believe their valuation for the first- and the third-party contents (i.e. v_f and v_t, respectively) to be uniformly and randomly distributed on the interval $[0, 1]$. Therefore, their optimal strategy is to compare the value they observe from the first content they sample to a reservation value. If the observed value is below the reservation value they sample the second content and if the observed value is above the reservation value, they stop sampling and consume the first content (Weitzman, 1979). Furthermore, we assume that the sampling a recommended content is costless to users. For example, when the platform recommends a content to a user, it provides valuable information (otherwise costly) about the content without any search costs. We note that this assumption simplifies the analysis but does not derive the results and can comfortably be relaxed. Finally, we assume that the search cost is small enough that users want to participate in the market.

Suppose the platform has full information with respect to a user's valuation for either content. Therefore, the platform may choose to provide a personalized recommendation to her. The user is aware of the platform's recommendation strategy and may choose to follow the recommendation and consume the recommended content without sampling the second one; or ignore the recommendation and sample the second content at a cost c. The rational user chooses the option that maximizes her utility. Given an exogenous royalty rate r, the platform pays a share of the net generated revenue that is proportional to the consumption of third-party content on the platform.

4 Analysis

4.1 Benchmark: Absence of Recommendation

We begin our analysis of the market in the absence of a recommendation system. In the absence of a recommendation system, a user resolves her valuation uncertainty upon sampling a content. Upon paying the subscription fee and joining the platform, the user randomly chooses one of the two contents to sample. Upon realizing her value, she may choose to consume the content or sample the alternative content at a cost c; in this case, she subsequently consumes the content with the highest realized value. First, we characterize the user's optimal search rule. Suppose a user realizes the value of the first content she samples to be $\tilde{v}_1$. Then her expected gain from sampling the alternative content is $\int_{\tilde{v}_1}^{1}(v_2 - \tilde{v}_1)f(v_2)dv_2 = \frac{1}{2} + \frac{1}{2}\tilde{v}_1^2 - \tilde{v}_1$. She samples the content if this expected gain is weakly greater than the search cost she must incur. Let us denote V as the reservation value of the user which is the solution to the following equality:

$$\frac{1}{2} + \frac{1}{2}V^2 - V = c \Leftrightarrow V = 1 - \sqrt{2c}$$

Therefore the optimal search policy is such that the user compares her realized value $\tilde{v}_1$ with her reservation value V: she samples the alternative content if $\tilde{v}_1 < V$; otherwise, she stops and consumes the first content immediately.

Given the optimal search policy, a user's expected utility from joining the platform is given by:

$$\begin{aligned} U(s) &= \int_{1-\sqrt{2c}}^{1} v_1\,dv_1 + \int_{0}^{1-\sqrt{2c}} -c\,dv_1 + \int_{0}^{1-\sqrt{2c}}\int_{v_1}^{1} v_2\,dv_2\,dv_1 + \int_{0}^{1-\sqrt{2c}}\int_{0}^{v_1} v_1\,dv_2\,dv_1 - s \\ &= \frac{2}{3} - c + \frac{1}{3}\sqrt{8c^3} - s \end{aligned}$$

The platform's optimal subscription fee is to extract all user's surplus (i.e. $s^* = \frac{2}{3} - c + \frac{1}{3}\sqrt{8c^3}$). The net generated revenue equals s^*. The probability that the third-party content is consumed (i.e. $v_f \geq v_t$) equals $\frac{1}{2}$. Therefore, the platform keeps $\left(\frac{2}{3} - c + \frac{1}{3}\sqrt{8c^3}\right)\left(1 - \frac{1}{2}r\right)$ and leaves the rest to the third-party provider.

4.2 Benchmark: Neutral Recommendation

Consider the benchmark case where there exists no recommendation bias. The platform recommends f to a user if her value for the first-party content is weakly greater than that of the third-party content (i.e. $v_f \geq v_t$) and t, otherwise. Users are rational and fully aware of the platform's recommendation strategy. Therefore, they consume the content that the platform recommends to them. A user's expected utility from joining the platform is given by:

$$U(s) = \int_{0}^{1}\int_{v_t}^{1} v_f\,dv_f\,dv_t + \int_{0}^{1}\int_{0}^{v_t} v_t\,dv_f\,dv_t - s = \frac{2}{3} - s \tag{4.1}$$

The platform's optimal subscription fee is to extract all user's surplus (i.e. $s^* = \frac{2}{3}$). Users join the service and consume the content that is recommended to them. The net generated revenue equals s^*. The probability that the first-party content is consumed (i.e. $v_f \geq v_t$) equals $\int_0^1 \int_{v_t}^1 1\,dv_f\,dv_t = \frac{1}{2}$. Subsequently, the probability that the third-party content is consumed equals $\frac{1}{2}$. Therefore, the platform keeps $\frac{2}{3}\left(1 - \frac{1}{2}r\right)$ and leaves the rest to the third-party provider. A neutral recommendation system relays the platform's private information to the users, enabling the platform to extract the whole surplus generated from the bundle.

4.3 Biased Recommendation

Suppose the platform adopts a recommendation strategy such that the platform recommends the first-party content to a user if $v_f \geq v_t + \beta$ (as depicted in Figure 4·1). Thus, for a $\beta > 0$, the platform recommends her own content to a user whose value for the third-party content is $v_t \leq v_f - \beta$. Suppose a user joins the service, receives a recommendation t, samples the content t and realizes her value $\tilde{v}_t$. Recall that we assume a user realizes her valuation of the recommended content without incurring any cost. She updates her belief about her valuation of the first-party content (v_f) such that $v_f^u \sim U[0, \tilde{v}_t - \beta]$. Note the upper limit of the user's valuation for the first-party content equals $\tilde{v}_t - \beta$ and is strictly smaller than $\tilde{v}_t$. Therefore, the user's does not expect any gain from sampling the first-party content, and follows the recommendation of the platform and consumes the third-party content.

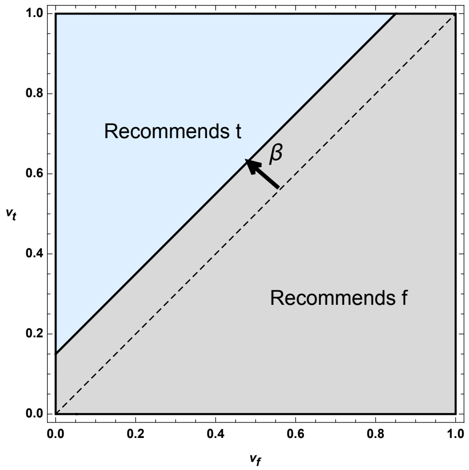

Figure 4·1: The platform's biased recommendation strategy

Now suppose a user who receives a recommendation f and samples the f content. She realizes her value to be $\tilde{v}_f$ and updates her belief about v_t such that:

- If $1 - \beta \leq \tilde{v}_f \leq 1$, $v_t^u \sim U[0, 1]$. Therefore, the platform's recommendation of the first-party content does not provide any additional information about her valuation of the third-party content.

- If $0 \leq \tilde{v}_f < 1 - \beta$, $v_t^u \sim U[0, \tilde{v}_f + \beta]$.

Consequently, a user who receives a recommendation f and her realized value of the first-party content is such that $\tilde{v}_f \geq 1 - \beta$, samples the third-party content if her realized value is less than or equal her reservation value, denoted by x. Note that x is the value for which the expected gain from sampling the third-party content equals the user's search cost, and is the solution to the equality:

$$\int_{x}^{1} (v_t - x)\, dv_t = c \Leftrightarrow x = 1 - \sqrt{2c} \tag{4.2}$$

Given her reservation value $x = 1 - \sqrt{2c}$, a user with a realized value $\tilde{v}_f$, where, $1 - \beta \leq \tilde{v}_f$, samples the second content if $\beta > \sqrt{2c}$ and $\tilde{v}_f < 1 - \sqrt{2c}$.

Meanwhile, a user who receives a recommendation f and her realized value of the first-party content is such that $0 \leq \tilde{v}_f < 1 - \beta$, samples the third-party content if her realized value is less than or equal her reservation value, denoted by X. Where X denotes the value for which the expected gain from sampling the third-party content equals the search cost and is the solution to the following equation:

$$\int_{X}^{X+\beta} (v_t - X)\, f(v_t)\, dv_t = \int_{X}^{X+\beta} \frac{v_t - X}{X + \beta}\, dv_t = c \Leftrightarrow X = \frac{\beta(\beta - 2c)}{2c} \tag{4.3}$$

Consequently, a user with $\tilde{v}_f < 1 - \beta$, searches the second content if $2c < \beta$ and $\tilde{v}_f < \frac{\beta(\beta-2c)}{2c}$.

When $\beta > \sqrt{2c}$, we have, $\frac{\beta(\beta-2c)}{2c} > 1 - \sqrt{2c}$. Therefore, a user who receives a recommendation f, samples the third-party content if her realized value is $\tilde{v}_f < 1 - \sqrt{2c}$. Intuitively, when the recommendation system is heavily biased toward the first-party content (i.e. $\beta > \sqrt{2c}$), the user finds it optimal to search the third-party content unless her realized value of the first-party content is $\tilde{v}_f \geq 1 - \sqrt{2c}$. When $\beta < \sqrt{2c}$, we have, $\frac{\beta(\beta-2c)}{2c} < 1 - \sqrt{2c}$. Therefore, if $2c < \beta \leq \sqrt{2c}$, and a user receives a recommendation f, she samples the third-party content if $\tilde{v}_f < \frac{\beta(\beta-2c)}{2c}$, and consumes the content with the highest realized value. Finally, if $\beta \leq 2c$, the user does not find it worthwhile to sample the third-party content at a cost c, regardless of her realized value $\tilde{v}_f$ and consumes the first-party content. Intuitively, when the recommendation system is only slightly biased toward the first-party content, the user's expected gain from sampling the third-party content does not compensate her for the search cost she must incur.

Let us first consider the users' expected utility from subscribing to the bundle for a fee s, when $\beta \leq 2c$.

$$\begin{aligned} U(s) &= \int_{\beta}^{1}\int_{v_t-\beta}^{1} v_f \, dv_f \, dv_t + \int_{0}^{\beta}\int_{0}^{1} v_f \, dv_f \, dv_t + \int_{\beta}^{1}\int_{0}^{v_t-\beta} v_t \, dv_f \, dv_t \\ &= \frac{2}{3} - \frac{1}{2}\beta^2 + \frac{1}{3}\beta^3 - s \end{aligned}$$

The first two expressions in the user's expected utility denotes the her expected value from consuming the first-party content when v_f and v_t are such that the platform recommends f (i.e. $v_f \geq \max(v_t - \beta, 0)$). The third expression in the user's expected utility denotes the user's expected value from consuming the third-party content if her v_f and v_t are such that the platform recommends t (i.e. $v_t > v_f + \beta$). Given the users' expected utility from subscribing to the bundle, the platform's optimal pricing strategy is to set $s^* = \frac{2}{3} - \frac{1}{2}\beta^2 + \frac{1}{3}\beta^3$ and extract all surplus. Note that for any value of $0 < \beta \leq 2c$, the price of the bundle is less than the benchmark with the neutral recommendation. This is because the user expects a lower utility from joining a service with a biased recommendation system. Finally, the share of consumption of the third-party content on the platform is given by:

$$\int_{\beta}^{1}\int_{0}^{v_t-\beta} v_t\,dv_f\,dv_t = \frac{1}{2}\left(1-\beta\right)^2$$

Therefore, the platform keeps $\left(\frac{2}{3}-\frac{1}{2}\beta^2+\frac{1}{3}\beta^3\right)\left(1-\frac{1}{2}(1-\beta)^2r\right)$ and leaves the rest to the third-party provider. Next, we consider the case where $2c<\beta<\sqrt{2c}$, in which case a user with a realized value $\tilde{v}_f \leq X$, where $X=\frac{\beta(\beta-2c)}{2c}$, samples the third-party content at a cost c, and consumes the content with the highest realized value. Given the users' optimal search strategy, we write the expected utility from subscribing to the bundle as follows:

$$\begin{aligned} U(s) &= \int_{X}^{1-\beta}\int_{0}^{v_f+\beta} v_f\,dv_t\,dv_f + \int_{1-\beta}^{1}\int_{0}^{1} v_f\,dv_t\,dv_f + \int_{0}^{X}\int_{0}^{v_f} v_f\,dv_t\,dv_f + \\ &\quad \int_{0}^{1-\beta}\int_{v_f+\beta}^{1} v_t\,dv_t\,dv_f + \int_{0}^{X}\int_{v_f}^{v_f+\beta} v_t\,dv_t\,dv_f + f - c\int_{0}^{X}\int_{0}^{v_f+\beta} 1\,dv_t\,dv_f - s \\ &= \frac{2}{3}-\frac{\beta^3}{6}-\frac{1}{2}\beta^2(1-c)+\frac{\beta^4}{8c}-s \end{aligned}$$

The first two expressions denote the expected value from consuming the first-party content, where v_f and v_t are such that the user receives a recommendation f and consumes the first-party content without sampling the third-party content. The third expression denotes the expected value given the user samples the third-party content, but goes back to consuming the first content. The forth expression denotes the expected value when v_f and v_t are such that the user receives a recommendation t and consumes the third-party content. While, the fifth expression denotes the expected value from consuming the third-party content, only if the user receives a recommendation f, but samples the third-party content and realized that her value for the third-party content is greater than that of the first-party content. Finally, the last expression denotes the probability that the user realizes a value v_f, upon getting a recommendation f such that her best course of action is to sample the third-party content for a cost c.

Given the user's expected utility from subscribing to the bundle, the platform's optimal subscription fee is such that $s^* = \frac{2}{3} - \frac{\beta^3}{6} - \frac{1}{2}\beta^2(1-c) + \frac{\beta^4}{8c}$ and all surplus is extracted. Moreover, the share of consumption of the third-party content is as follows:

$$\int\limits_{0}^{1-\beta}\int\limits_{v_f+\beta}^{1} 1\,dv_t\,dv_f + \int\limits_{0}^{X}\int\limits_{v_f}^{v_f+\beta} 1\,dv_t\,dv_f = \frac{1}{2}(1-\beta)^2 + \frac{\beta^2(\beta-2c)}{2c}$$

Given the platform's optimal subscription fee, $s^* = \frac{2}{3} + \beta c - \beta^2 - \frac{\beta^3}{6} + \frac{\beta^4}{4c}$, the platform keeps a share equal to $\left(1 - \left(\frac{1}{2}(1-\beta)^2 + \frac{\beta^2(\beta-2c)}{2c}\right) r\right)$ and leaves the rest to the third-party provider.

Finally, if β is such that $\beta > \sqrt{2c}$, a user who receives a recommendation to consume the first-party content, and $\tilde{v}_f < 1 - \sqrt{2c}$, samples the third-party content. Therefore, the user's expected utility from subscribing to the bundle is such that:

$$\begin{aligned} U(s) \;=\;& \int\limits_{x}^{1}\int\limits_{0}^{1} v_f\,dv_t\,dv_f + \int\limits_{0}^{x}\int\limits_{0}^{v_f} v_f\,dv_t\,dv_f + \int\limits_{0}^{1-\beta}\int\limits_{v_f+\beta}^{1} v_t\,dv_t\,dv_f + \int\limits_{0}^{1-\beta}\int\limits_{v_f}^{v_f+\beta} v_t\,dv_t\,dv_f + \\ & \int\limits_{1-\beta}^{x}\int\limits_{v_f}^{1} v_t\,dv_t\,dv_f - c\left(\int\limits_{0}^{1-\beta}\int\limits_{0}^{v_f+\beta} 1\,dv_t\,dv_f + \int\limits_{1-\beta}^{x}\int\limits_{0}^{1} 1\,dv_t\,dv_f\right) - s \\ =\;& \frac{2}{3} - \frac{1}{2}c + \frac{2}{3}\sqrt{8c^3} + \frac{\beta^2 c}{2} - \beta c - s \end{aligned}$$

The first expression denotes the users' expected value from consuming the first-party content, when v_f and v_t are such that the user receives a recommendation f, and consume the first-party content, immediately. While, the second expression denotes the expected value from consuming the first-party content, when a user receives a recommendation f, and after realizing her value to be less than x, chooses to sample the third-party content. The third expression denote the expected value from consuming the third-party content, given the user receives a recommendation t. While, the fourth and fifth expressions, denote the expected value from consuming the third-party content when the user receives a recommendation f, samples the third-party content and realized her value for the third-party content to be greater than that of the first-party content. Finally, the last two expressions denote the probability that the user realizes a value v_f, upon getting a recommendation f such that her best course of action is to sample the third-party content for a cost c.

Given the user's expected utility from subscribing to the bundle, the platform's optimal subscription fee is such that $s^* = \frac{2}{3} - \frac{1}{2}c + \frac{2}{3}\sqrt{8c^3} + \frac{\beta^2 c}{2} - \beta c$ and all surplus is extracted. Moreover, the share of consumption of the third-party content is as follows:

$$\int_{\beta}^{1}\int_{0}^{v_t-\beta} dv_f dv_t + \int_{1-\beta}^{1-\sqrt{2c}}\int_{v_f}^{1} dv_t dv_f + \int_{0}^{1-\beta}\int_{v_f}^{v_f+\beta} dv_t dv_f = \frac{1}{2} - c$$

The platform chooses β^* that maximizes the expected revenue, given the users' optimal search policy and the optimal subscription fee $(s^*(\beta))$, induced by the platform's recommendation policy as follows:

$$\max_{\beta} s^*(\beta)(1 - r\kappa(\beta))$$

Where $\kappa(\beta)$ denotes the expected consumption rate of the third-party content given β.

Proposition 4.1. *A biased recommendation system recommends the first-party content to a user if $v_f \geq v_t - \beta^*$ and, the the second-party content, otherwise. For a search cost c there exists a unique threshold (c_{th}) such that:*

i If $c \leq c_{th}$, when $r \leq \frac{6c(1-2c)}{2-5c-24c^2+68c^3-56c^4}$, β^* satisfies $\beta^* \leq 2c$, and each user *consumes the recommended content. When* $r > \frac{6c(1-2c)}{2-5c-24c^2+68c^3-56c^4}$, β^* *satisfies* $2c < \beta^* < \sqrt{2c}$, *and a user who receives an* f *recommendation samples the third-party content if her realized value is* $\tilde{v}_f < \frac{\beta^*(\beta^*-2c)}{2c}$.

ii If $c > c_{th}$, β^* *satisfies* $\beta^* \leq 2c$, *and each user consumes the recommended content.*

Proposition 4.1 suggests that the platform's optimal bias toward first-party content depends on the magnitude of the users' search cost, and the third-party provider's royalty rate. The result is illustrated in Figure 4·2. Observe that given a small search cost c, the optimal β^* increases in r. The platform initially opts to set β^* such that the users take the platform's recommendation and do not sample the alternative. As r increases, the platform gradually shifts its strategy toward allowing some users to sample the third-party content, given a recommendation f.

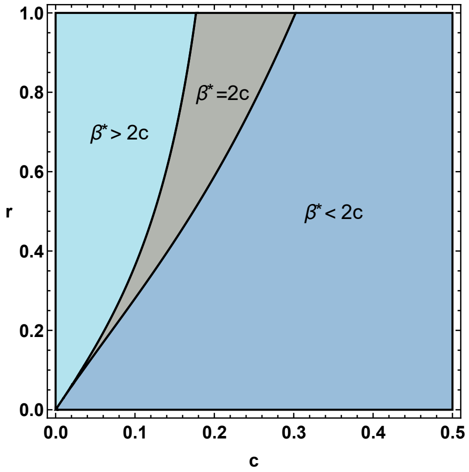

Figure 4·2: The platform's Optimal recommendation strategy as a function of c and r

Proposition 4.2 establishes that the price of the bundle and therefore, the total social welfare generated from the bundle, is decreasing in c and r. The proposition suggests that when search is costly, a biased recommendation may induce some users to sample the alternative content at a high cost. Therefore, the value of the bundle decreases for the users, forcing the platform to charge lower prices.

Proposition 4.2. *The price of the bundle is non-monotonically decreasing in the user's search cost and the third-party provider's royalty rate.*

We now consider how the design of a biased personalized recommendation service affects the platform's and the third-party provider's expected profit as compared to the benchmarks. With the adoption of biased recommendation system, the platform is able to steer consumption to increase its own revenue at the cost of lower third-party provider's revenue and social welfare. As the royalty rate of the third-party provider increases, the platform becomes more aggressive in steering users to consuming first-party content. However, the strategy is most effective when users incur high costs to search for their optimal content. This result is illustrated in Figure 4·3.

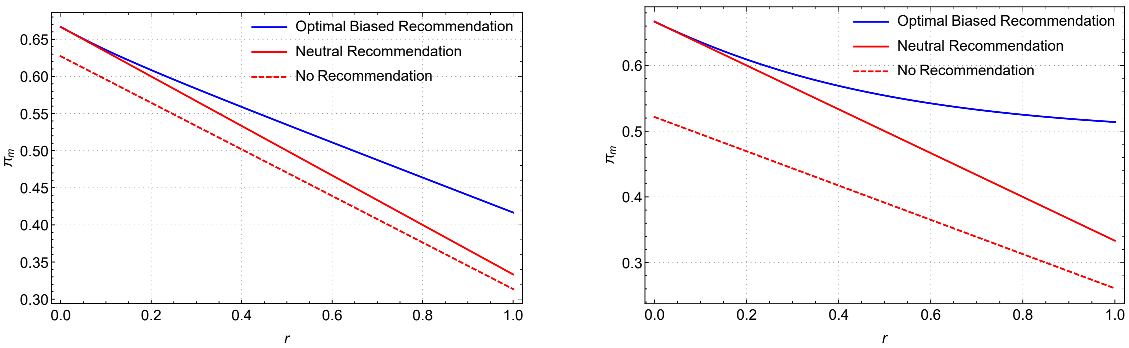

(a) $c = 0.05$ **(b)** $c = 0.3$

Figure 4·3: Platform's optimal revenue parameter a function of the third-party provider's royalty rate.

As illustrated in Figure 4·4, when search cost is sufficiently small, the revenue of the third-party provider is strictly increasing in r. Intuitively, a small search cost hinders the platform's attempt to steer users to first-party content. Namely, a strategic user with small search cost samples the third-party content, even if the platform recommends the first-party content to her. Therefore, the provider is strictly better off with higher royalty rates to extract its share of the revenue. However, when the search cost is high, it becomes easier for the platform to steer users to the first-party content. Taking into account that the platform's bias is increasing in r and that the price of the bundle is decreasing in r, the third-party provider's revenue is first increasing and then decreasing in r. This implies that the third-party provider may be better of foregoing a higher royalty rate for a higher share of consumption.

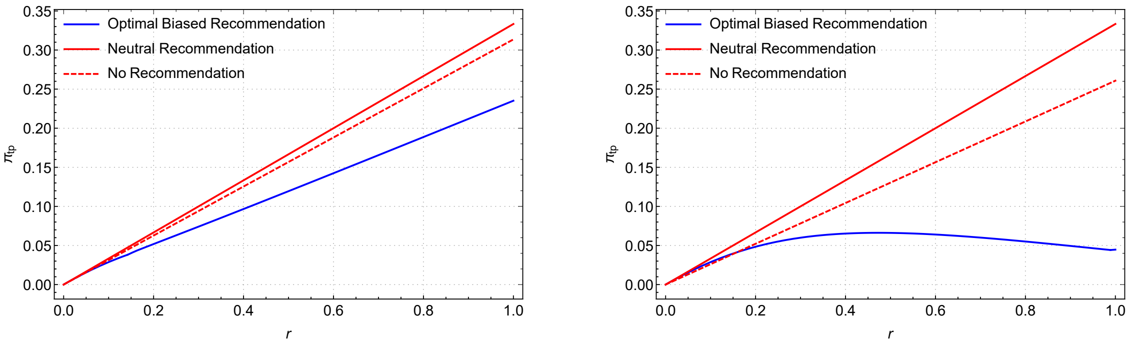

(a) $c = 0.05$ **(b)** $c = 0.3$

Figure 4·4: Revenue of the third-party provider as a function of the third-party provider's royalty rate.

5 Discussion and Conclusion

This paper studies the optimal design of the recommendation system of an streaming platform that has private information about users valuations and provides users with personalized recommendations on what content to consume. We show that the platform may utilize its recommendation system to steer users to consuming the content provided by the platform. Users take the recommendation of the platform and choose whether to consume the content or continue sampling to find the content that best matches their tastes. We show a linearly-biased recommendation system allows the platform to increase its revenue, albeit at the cost of lower social welfare.

When the search cost is sufficiently small, under the optimal recommendation strategy, some users may choose to ignore the recommendation of the platform and sample the alternative content. However, as the search cost increases the platform is able to steer users toward first-party content more aggressively. Furthermore, we show that when the search cost is high, the third-party provider is better off with an intermediate royalty rate.

References

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

Armstrong, M. and Zhou, J. (2011). Paying for prominence. *The Economic Journal*, 121(556):F368–F395.

Bimpikis, K., Papanastasiou, Y., and Zhang, W. (2020). Information provision in two-sided platforms: Optimizing for supply. *Available at SSRN 3617351*.

Bourreau, M. and Gaudin, G. (2018). Streaming platform and strategic recommendation bias. *Available at SSRN: https://ssrn.com/abstract=3290617*.

Cachon, G. P. and Feldman, P. (2011). Pricing services subject to congestion: Charge per-use fees or sell subscriptions? *Manufacturing & Service Operations Management*, 13(2):244–260.

Choi, H. and Mela, C. F. (2019). Monetizing online marketplaces. *Marketing Science*, 38(6):948–972.

Choudhary, V. and Zhang, Z. (2019). Product recommendation and consumer search. *Available at SSRN 3398994*.

De Corniere, A. and Taylor, G. (2019). A model of biased intermediation. *The RAND Journal of Economics*, 50(4):854–882.

Dinerstein, M., Einav, L., Levin, J., and Sundaresan, N. (2018). Consumer price search and platform design in internet commerce. *American Economic Review*, 108(7):1820–59.

Feldman, P., Frazelle, A. E., and Swinney, R. (2019). Can delivery platforms benefit restaurants? *Available at SSRN 3258739*.

Gaur, V. and Liu, X. (2020). Personalized recommendation system design for an online b2b platform. *Available at SSRN 3902710*.

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

Hagiu, A., Teh, T.-H., and Wright, J. (2020). Should platforms be allowed to sell on their own marketplaces? *Available at SSRN: https://ssrn.com/abstract=3606055*.

Hagiu, A. and Wright, J. (2015). Marketplace or reseller? *Management Science*, 61(1):184–203.

Johnson, J. P. (2017). The agency model and mfn clauses. *The Review of Economic Studies*, 84(3):1151–1185.

Lei, Z. and Swinney, R. (2018). Revenue sharing and subscription platforms for digital content distribution. *Available at SSRN 3314345*.

Papanastasiou, Y., Bimpikis, K., and Savva, N. (2018). Crowdsourcing exploration. *Management Science*, 64(4):1727–1746.

Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, pages 641–654.

Zhou, B. and Zou, T. (2021). Competing for recommendations: The strategic impact of personalized product recommendations in online marketplaces. *Available at SSRN 3761350*.

Appendix

Proof of Proposition 4.1

Proof. First, let us consider the platform's maximization problem where $\sqrt{2c} \leq \beta < 1$ and a user samples the third-party content if she receives a recommendation f and realizes her value of the first-party content to be less than her reservation value $1 - \sqrt{2c}$.

$$\max_{\beta} \left(\frac{2}{3} - \frac{1}{2}c + \frac{2}{3}\sqrt{8c^3} + \frac{\beta^2 c}{2} - \beta c\right)\left(1 - \left(\frac{1}{2} - c\right) r\right)$$
$$s.t. \quad \sqrt{2c} \leq \beta < 1$$

The revenue function is strictly decreasing in β (ie. $\frac{\partial \pi}{\partial \beta} = -c\left(1 - \left(\frac{1}{2} - c\right) r\right)(1-\beta) < 0$). Therefore, when $\sqrt{2c} \leq \beta < 1$, $\beta^* = \sqrt{2c}$. Next, we consider the platform's maximization problem given $2c \leq \beta \leq \sqrt{2c}$, as follows:

$$\max_{\beta} \left(\frac{2}{3} - \frac{\beta^3}{6} - \frac{1}{2}\beta^2(1-c) + \frac{\beta^4}{8c}\right)\left(1 - \left(\frac{1}{2}(1-\beta)^2 + \frac{\beta^2(\beta - 2c)}{2c}\right) r\right)$$
$$s.t. \quad 2c < \beta \leq \sqrt{2c}$$

Solving the first-order condition, we have:

$$\begin{aligned}\frac{\partial \pi}{\partial \beta} \quad = \quad & -\left(\frac{2}{3} - \frac{\beta^3}{6} - \frac{1}{2}\beta^2(1-c) + \frac{\beta^4}{8c}\right)\left(-1 + \beta + \frac{\beta^2}{2c} + \frac{\beta(\beta-2c)}{c}\right) r - \\ & \left(-\frac{\beta^2}{2} - \beta(1-c) + \frac{\beta^3}{2c}\right)\left(\frac{1}{2}(1-\beta)^2 + \frac{\beta^2(\beta-2c)}{2c}\right) r - \frac{\beta^2}{2} - \beta(1-c) + \frac{\beta^3}{2c}\end{aligned}$$

If $\frac{6c(1-2c)}{2-5c-24c^2+68c^3-56c^4} < r \leq 1$, the second-order condition is satisfied, and the the revenue function has a local maximum at β^*, where, β^* is the implicit solution to the first-order condition. However, if $r < \frac{6c(1-2c)}{2-5c-24c^2+68c^3-56c^4}$, the function is strictly decreasing in β and $\beta^* = 2c$.

Finally, we have:

$$\max_{\beta} \left(\frac{2}{3} - \frac{1}{2}\beta^2 + \frac{1}{3}\beta^3\right)\left(1 - \frac{1}{2}(1-\beta)^2 r\right)$$
$$s.t. \quad 0 \leq \beta \leq 2c$$

The first-order condition gives:

$$\frac{\partial \pi}{\partial \beta} = (1-\beta)\left(\frac{2}{3}-\frac{\beta^2}{2}+\frac{\beta^3}{3}\right)r-\beta\,(1-\beta)\left(1-\frac{1}{2}(1-\beta)^2 r\right)=0$$

If $r < \frac{6c}{2+3c-18c^2+20c^3}$, the function has a local maximum which is the implicit solution to the first-order condition and that satisfies the second-order condition. However, if $r \geq \frac{6c}{2+3c-18c^2+20c^3}$, the function is strictly increasing in β and $\beta^* = 2c$. Comparing the optimal revenue function given the users' optimal search behavior, we have:

- If $r < \frac{6c}{2+3c-18c^2+20c^3}$, β^* is the implicit solution to the equation:

$$\frac{\partial}{\partial \beta}\left(\left(\frac{2}{3}-\frac{1}{2}\beta^2+\frac{1}{3}\beta^3\right)\left(1-\frac{1}{2}(1-\beta)^2 r\right)\right)=0$$

- If $\frac{6c}{2+3c-18c^2+20c^3} \leq r \leq \frac{6c(1-2c)}{2-5c-24c^2+68c^3-56c^4}$, $\beta^* = 2c$.
- If $r > \frac{6c(1-2c)}{2-5c-24c^2+68c^3-56c^4}$, β^* is the implicit solution to the equation:

$$\frac{\partial}{\partial \beta}\left(\left(\frac{2}{3}-\frac{\beta^3}{6}-\frac{1}{2}\beta^2(1-c)+\frac{\beta^4}{8c}\right)\left(1-\left(\frac{1}{2}(1-\beta)^2+\frac{\beta^2(\beta-2c)}{2c}\right)r\right)\right)=0$$

When c is such that $\frac{6c}{2+3c-18c^2+20c^3} > 1$, given the optimal recommendation strategy described above, β^* is the implicit solution to the equation:

$$\frac{\partial}{\partial \beta}\left(\left(\frac{2}{3}-\frac{1}{2}\beta^2+\frac{1}{3}\beta^3\right)\left(1-\frac{1}{2}(1-\beta)^2 r\right)\right)=0$$

(i.e. the second root of $4r-3(2-r)\beta-9r\beta^2+5r\beta^3=0$), and does not depend on the user's search cost c. If c is such that $\frac{6c}{2+3c-18c^2+20c^3} \leq 1$ and $\frac{6c(1-2c)}{2-5c-24c^2+68c^3-56c^4} > 1$, for $r < \frac{6c}{2+3c-18c^2+20c^3}$, β^* is independent of c. However, if $r \geq \frac{6c}{2+3c-18c^2+20c^3}$, $\beta^* = 2c$, is strictly increasing in c. Finally, if c is such that $\frac{6c(1-2c)}{2-5c-24c^2+68c^3-56c^4} < 1$, we have: 1) for $r < \frac{6c}{2+3c-18c^2+20c^3}$, β^* is independent of c; 2) for $\frac{6c}{2+3c-18c^2+20c^3} \leq r \leq \frac{6c(1-2c)}{2-5c-24c^2+68c^3-56c^4}$, $\beta^* = 2c$, is strictly increasing in c; 3) for $\frac{6c(1-2c)}{2-5c-24c^2+68c^3-56c^4} < r \leq 1$, β^* is the implicit solution to the equation:

$$\frac{\partial}{\partial \beta}\left(\left(\frac{2}{3}-\frac{\beta^3}{6}-\frac{1}{2}\beta^2(1-c)+\frac{\beta^4}{8c}\right)\left(1-\left(\frac{1}{2}(1-\beta)^2+\frac{\beta^2(\beta-2c)}{2c}\right)r\right)\right)=0$$

, and is strictly increasing in c. □

Proof of Proposition 4.2

Proof. In the proof of Proposition 4.1, we establish that β^*, is non-monotonically increasing in c. When $r < \frac{6c}{2+3c-18c^2+20c^3}$, β^* and consequently the price of the bundle, given by $\frac{2}{3} - \frac{1}{2}\beta^2 + \frac{1}{3}\beta^3$, does not depend on c. However, β^* is strictly increasing in r and the price is strictly decreasing in β. Therefore, the price of the bundle is decreasing in r.

When $\frac{6c(1-2c)}{2-5c-24c^2+68c^3-56c^4} < 1$, for $\frac{6c(1-2c)}{2-5c-24c^2+68c^3-56c^4} < r \leq 1$, β^* is strictly increasing in c (refer to proof of Proposition 4.1). While, the price of the bundle given by: $s = \left(\frac{2}{3} - \frac{\beta^3}{6} - \frac{1}{2}\beta^2(1-c) + \frac{\beta^4}{8c}\right)$ at β^* is decreasing in c. Note with the envelop theorem we have:

$$\frac{\partial s}{\partial c} = \frac{\partial s}{\partial c} + \frac{\partial s}{\partial \beta}\frac{\partial \beta}{\partial c} < 0$$

where $= \frac{\partial s}{\partial c} = \frac{1}{8}\beta^{*2}\left(4 - \frac{\beta^{*2}}{c^2}\right) < 0$, $\frac{\partial s}{\partial \beta} = \frac{1}{2}\beta^{*2} - (1-c)\beta^* + \frac{\beta^{*3}}{2c} < 0$ and $\frac{\partial \beta}{\partial c} > 0$.

Moreover, given that $\frac{\partial s}{\partial r} = \frac{\partial s}{\partial r} + \frac{\partial s}{\partial \beta}\frac{\partial \beta}{\partial r} = \frac{\partial s}{\partial \beta}\frac{\partial \beta}{\partial r}$, and β^* is increasing in r, the price of the bundle is decreasing in r.

Finally, when $\frac{6c}{2+3c-18c^2+20c^3} \leq r \leq \frac{6c(1-2c)}{2-5c-24c^2+68c^3-56c^4}$ and $\beta^* = 2c$, the price of the bundle is decreasing in c, but independent of r. □

References

Abdallah, T., Asadpour, A., and Reed, J. (2021). Large-scale bundle-size pricing: A theoretical analysis. *Operations Research*.

Alaei, S., Makhdoumi, A., and Malekian, A. (2019). Optimal subscription planning for digital goods. *Available at SSRN 3476296*.

Altug, M. S. and Aydinliyim, T. (2016). Counteracting strategic purchase deferrals: The impact of online retailers' return policy decisions. *Manufacturing & Service Operations Management*, 18(3):376–392.

Anderson, C. (2006). *The long tail: Why the future of business is selling less of more*. Hachette Books.

Anderson Jr, E. G., Parker, G. G., and Tan, B. (2014). Platform performance investment in the presence of network externalities. *Information Systems Research*, 25(1):152–172.

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

Armstrong, M. and Weeds, H. (2007). Programme quality in subscription and advertising-funded television. *Unpublished working paper*.

Armstrong, M. and Wright, J. (2007). Two-sided markets, competitive bottlenecks and exclusive contracts. *Economic Theory*, 32(2):353–380.

Armstrong, M. and Zhou, J. (2011). Paying for prominence. *The Economic Journal*, 121(556):F368–F395.

Aviv, Y. and Pazgal, A. (2008). Optimal pricing of seasonal products in the presence of forward-looking consumers. *Manufacturing & service operations management*, 10(3):339–359.

Aviv, Y. and Vulcano, G. (2012). Dynamic list pricing. In *The Oxford handbook of pricing management*.

Aydin, G. and Ziya, S. (2009). Personalized dynamic pricing of limited inventories. *Operations Research*, 57(6):1523–1531.

Bakos, Y. and Brynjolfsson, E. (1999). Bundling information goods: Pricing, profits, and efficiency. *Management science*, 45(12):1613–1630.

Ban, G.-Y. and Keskin, N. B. (2020). Personalized dynamic pricing with machine learning: High dimensional features and heterogeneous elasticity. *Available at SSRN 2972985*.

Bernstein, F., Chakraborty, S., and Swinney, R. (2020). Dynamic content variation with customer learning. *Available at SSRN 3489425*.

Bernstein, F., Kök, A. G., and Xie, L. (2015). Dynamic assortment customization with limited inventories. *Manufacturing & Service Operations Management*, 17(4):538–553.

Bernstein, F., Modaresi, S., and Sauré, D. (2019). A dynamic clustering approach to data-driven assortment personalization. *Management Science*, 65(5):2095–2115.

Bhargava, H. K. (2021). Bundling for flexibility and variety: An economic model for multiproducer value aggregation. *Management Science*, 67(4):2365–2380.

Bimpikis, K., Papanastasiou, Y., and Zhang, W. (2020). Information provision in two-sided platforms: Optimizing for supply. *Available at SSRN 3617351*.

Bourreau, M. and Gaudin, G. (2018). Streaming platform and strategic recommendation bias. *Available at SSRN: https://ssrn.com/abstract=3290617*.

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

Cachon, G. P. and Feldman, P. (2011). Pricing services subject to congestion: Charge per-use fees or sell subscriptions? *Manufacturing & Service Operations Management*, 13(2):244–260.

Cachon, G. P. and Feldman, P. (2018). Pricing capacity over time and recourse strategies: Facilitate reselling, offer refunds/options, or overbook? *Available at http://dx.doi.org/10.2139/ssrn/3235757*.

Cachon, G. P. and Swinney, R. (2009). Purchasing, pricing, and quick response in the presence of strategic consumers. *Management Science*, 55(3):497–511.

Caillaud, B. and Jullien, B. (2003). Chicken & egg: Competition among intermediation service providers. *RAND journal of Economics*, pages 309–328.

Calvano, E. and Polo, M. (2020). Strategic differentiation by business models: Free-to-air and pay-tv. *The Economic Journal*, 130(625):50–64.

Candogan, O., Bimpikis, K., and Ozdaglar, A. (2012). Optimal pricing in networks with externalities. *Operations Research*, 60(4):883–905.

Caro, F. and Gallien, J. (2007). Dynamic assortment with demand learning for seasonal consumer goods. *Management Science*, 53(2):276–292.

Carroni, E., Madio, L., and Shekhar, S. (2020). Superstars in two-sided markets: exclusives or not? *Available at SSRN 3243777*.

Chen, T., Fenyo, K., Yang, S., and Zhang, J. (2018). Thinking inside the subscription box: New research on e-commerce consumers. *McKinsey & Company*, pages 1–9.

Cheung, W. C. and Simchi-Levi, D. (2017). Thompson sampling for online personalized assortment optimization problems with multinomial logit choice models. *Available at SSRN 3075658*.

Choi, H. and Mela, C. F. (2019). Monetizing online marketplaces. *Marketing Science*, 38(6):948–972.

Choudhary, V. and Zhang, Z. (2019). Product recommendation and consumer search. *Available at SSRN 3398994*.

De Corniere, A. and Taylor, G. (2019). A model of biased intermediation. *The RAND Journal of Economics*, 50(4):854–882.

DeGroot, M. H. (2005). *Optimal statistical decisions*, volume 82. John Wiley & Sons.

den Boer, A. V. (2015). Dynamic pricing and learning: historical origins, current research, and new directions. *Surveys in operations research and management science*, 20(1):1–18.

DeValve, L. and Pekeč, S. (2016). Optimal price/advertising menus for two-sided media platforms. *Available at SSRN 2715531*.

Dinerstein, M., Einav, L., Levin, J., and Sundaresan, N. (2018). Consumer price search and platform design in internet commerce. *American Economic Review*, 108(7):1820–59.

Feldman, P., Frazelle, A. E., and Swinney, R. (2019a). Can delivery platforms benefit restaurants? *Available at SSRN 3258739*.

Feldman, P., Papanastasiou, Y., and Segev, E. (2019b). Social learning and the design of new experience goods. *Management Science*, 65(4):1502–1519.

Forbes (2019). Why subscription boxes are here to stay.

Fudenberg, D. and Villas-Boas, J. M. (2006). Behavior-based price discrimination and customer recognition. *Handbook on economics and information systems*, 1:377–436.

Gaur, V. and Liu, X. (2020). Personalized recommendation system design for an online b2b platform. *Available at SSRN 3902710.*

Geng, X., Stinchcombe, M. B., and Whinston, A. B. (2005). Bundling information goods of decreasing value. *Management science*, 51(4):662–667.

Grand View Research (2021). *Video Streaming Market Size, Share & Trends Analysis Report By Streaming Type, By Solution, By Platform, By Service, By Revenue Model, By Deployment Type, By User, By Region, And Segment Forecasts, 2021 - 2028.* Grand View Research.

Guardian, T. (2018). Taylor swift strikes a blow for fellow artists as digital revenues soar. *The Guardian.*

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

Hagiu, A. and Lee, R. S. (2011). Exclusivity and control. *Journal of Economics & Management Strategy*, 20(3):679–708.

Hagiu, A. and Spulber, D. (2013). First-party content and coordination in two-sided markets. *Management Science*, 59(4):933–949.

Hagiu, A., Teh, T.-H., and Wright, J. (2020). Should platforms be allowed to sell on their own marketplaces? *Available at SSRN: https://ssrn.com/abstract=3606055.*

Hagiu, A. and Wright, J. (2015). Marketplace or reseller? *Management Science*, 61(1):184–203.

Hotelling, H. (1990). Stability in competition. In *The collected economics articles of Harold Hotelling*, pages 50–63. Springer.

Hu, X., Wan, Z., and Murthy, N. N. (2019). Dynamic pricing of limited inventories with product returns. *Manufacturing & Service Operations Management*, 21(3):501–518.

Ishihara, A. and Oki, R. (2020). Exclusive content in two-sided markets. *Available at SSRN 2941972.*

Jiang, H., Shilin, Z., and Zhengang, S. (2017). Two-sided market pricing in operations management: Review of current literature and research directions. *Management Science and Engineering*, 11(4):30–35.

Johnson, J. P. (2017). The agency model and mfn clauses. *The Review of Economic Studies*, 84(3):1151–1185.

Johnson, J. P. and Myatt, D. P. (2006). On the simple economics of advertising, marketing, and product design. *American Economic Review*, 96(3):756–784.

Kallus, N. and Udell, M. (2020). Dynamic assortment personalization in high dimensions. *Operations Research*.

Kök, A. G., Fisher, M. L., and Vaidyanathan, R. (2015). Assortment planning: Review of literature and industry practice. In *Retail supply chain management*, pages 175–236. Springer.

Lei, Z. and Swinney, R. (2018). Revenue sharing and subscription platforms for digital content distribution. *Available at SSRN 3314345*.

Li, K. J. (2021). Behavior-based quality discrimination. *Manufacturing & Service Operations Management*, 23(2):425–436.

Li, S., Luo, Q., Qiu, L., and Bandyopadhyay, S. (2020). Optimal pricing model of digital music: Subscription, ownership or mixed? *Production and Operations Management*, 29(3):688–704.

Mišić, V. V. and Perakis, G. (2020). Data analytics in operations management: A review. *Manufacturing & Service Operations Management*, 22(1):158–169.

Montgomery, A. L. and Smith, M. D. (2009). Prospects for personalization on the internet. *Journal of Interactive Marketing*, 23(2):130–137.

Murthi, B. and Sarkar, S. (2003). The role of the management sciences in research on personalization. *Management Science*, 49(10):1344–1362.

Papanastasiou, Y., Bimpikis, K., and Savva, N. (2018). Crowdsourcing exploration. *Management Science*, 64(4):1727–1746.

Papanastasiou, Y. and Savva, N. (2017). Dynamic pricing in the presence of social learning and strategic consumers. *Management Science*, 63(4):919–939.

Parker, G. G. and Van Alstyne, M. W. (2005). Two-sided network effects: A theory of information product design. *Management science*, 51(10):1494–1504.

Rhee, K.-E. and Thomadsen, R. (2017). Behavior-based pricing in vertically differentiated industries. *Management Science*, 63(8):2729–2740.

Rochet, J.-C. and Tirole, J. (2006). Two-sided markets: a progress report. *The RAND journal of economics*, 37(3):645–667.

Rusmevichientong, P., Shen, Z.-J. M., and Shmoys, D. B. (2010). Dynamic assortment optimization with a multinomial logit choice model and capacity constraint. *Operations research*, 58(6):1666–1680.

Sanchez-Cartas, J. M. and Leon, G. (2019). Multisided platforms and markets: A literature review. *SSRN Electron. J*, pages 1–62.

Shen, Q. and Miguel Villas-Boas, J. (2018). Behavior-based advertising. *Management Science*, 64(5):2047–2064.

Strauss, A. K., Klein, R., and Steinhardt, C. (2018). A review of choice-based revenue management: Theory and methods. *European Journal of Operational Research*, 271(2):375–387.

Su, X. (2010). Intertemporal pricing and consumer stockpiling. *Operations research*, 58(4-part-2):1133–1147.

Su, X. and Zhang, F. (2008). Strategic customer behavior, commitment, and supply chain performance. *Management Science*, 54(10):1759–1773.

Tan, B., Anderson Jr, E. G., and Parker, G. G. (2020). Platform pricing and investment to drive third-party value creation in two-sided networks. *Information Systems Research*, 31(1):217–239.

Tan, T. F., Netessine, S., and Hitt, L. (2017). Is tom cruise threatened? an empirical study of the impact of product variety on demand concentration. *Information Systems Research*, 28(3):643–660.

Times, T. N. Y. (2018). Netflix will keep ‘friends’ through next year in a $100 million agreement. *The New York Times*.

Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, pages 641–654.

Weyl, E. G. (2010). A price theory of multi-sided platforms. *American Economic Review*, 100(4):1642–72.

Xu, A., You, W., Zhong, X., and He, Q.-C. (2020). Strategic customer recommendations in online service platform. *Available at SSRN*.

Zennyo, Y. (2020). Platform encroachment and own-content bias. *Available at SSRN 3683287*.

Zhou, B. and Zou, T. (2021). Competing for recommendations: The strategic impact of personalized product recommendations in online marketplaces. *Available at SSRN 3761350*.

[redacted]

[redacted]

[redacted]

[redacted]