Carnegie Mellon University
Tepper School of Business

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of

DOCTOR OF PHILOSOPHY
INDUSTRIAL ADMINISTRATION
(MARKETING)

Titled

"Essays on Technology-Driven Marketing"

Presented by

Zijun (June) Shi

Accepted by

Kannan Srinivasan **4/19/2019**

Chair: Prof. Kannan Srinivasan **Date**

Approved by The Dean

Robert M. Dammon **4/28/2019**

Dean Robert M. Dammon **Date**

ESSAYS ON TECHNOLOGY-DRIVEN MARKETING

Zijun (June) Shi

Submitted to the Tepper School of Business in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Dissertation Committee

Kannan Srinivasan (Chair)

Alan Montgomery

Kaifu Zhang

Xiao Liu

Hui Li

May, 2019

Abstract:

With the development of technology in business applications, new marketing problems emerge, creating challenges for both practitioners and researchers. In this dissertation, I investigate marketing issues that involve new technology or require research methodologies enabled by new technology. I take an interdisciplinary approach, combining structural modeling, analytical modeling, machine learning, and causal inference, to study problems on pricing, media hype, and branding in three essays.

In the first essay, I examine the optimality of the freemium pricing strategy. Despite its immense popularity, the freemium business model remains a complex strategy to master and often a topic of heated debate. Adopting a generalized version of the screening framework à la Mussa and Rosen (1978), we ask when and why a firm should endogenously offer a zero price on its low-end product when users' product usages generate network externalities on each other. Our analysis indicates freemium can only emerge if the high- and low-end products provide asymmetric marginal network effects. In other words, the firm would set a zero price for its low-end product only if the high-end product provided larger utility gain from an expansion of the firm's user base. In contrast to conventional beliefs, a firm pursuing the freemium strategy might increase the baseline quality on its low-end product above the “efficient” level, which seemingly reduces differentiation.

In the second essay, I study how hype news from celebrity doctors affects the supply of information for weight-loss products. Consumers' healthcare choices are heavily influenced by public information. A distressing trend is desceptive information being propelled to popularity by trusted spokespeople. For example, Dr. Oz, a celebrity doctor, has made medical

recommendations directly against scientific evidence. Whether public information from reputable sources could correct misleading health information or not remains unknown. This study fills this research gap. By analyzing textual content using deep learning, I find that legitimate news outlets responded to *The Dr. Oz Show* by generating more news articles and carrying higher sentiment, hence amplifying rather than correcting hype news. Research articles reacted too slowly. Consumer reviews provided some correction but were overwhelmed by the opposite voice. Our findings have public policy implications on media content intervention and consumer protection.

In the third essay, I develop a dynamic structural model of fashion choices of brands and styles to investigate the implication of prohibiting fast fashion copycats, leveraging user-generated data from fashion-specific social media and deep learning methods on image analytics. I find that copycats can enhance high-end brands demand, contrary to conventional wisdom, due to several novel mechanisms: first, the affordability of mixing low-end copycats with high-end brands boosts demand for high-end brands from financially constrained consumers; second, good styles from low-end brands can help a consumer to build up his popularity/likeability, which increases his value for high-end brands and reduces the cost. Substantively, our results shed light on copyright enforcement and have implications on how fashion brands should react to copycats. Methodically, I developed a framework to analyze consumer choices where visual features are important product attributes and peer feedback hugely affects the decision-making process.

Acknowledgment:

I am deeply indebted to the support and guidance of my dissertation committee members: Professor Kannan Srinivasan, Professor Kaifu Zhang, Professor Xiao Liu, Professor Alan Montgomery, and Professor Hui Li. Since the beginning of my five-year journey, I was constantly enlightened by the discussion with Professor Kannan Srinivasan. Throughout the writing of this dissertation, I have received a great deal of support and utmost encouragement from him. Professor Kaifu Zhang guided me on analytical research in Marketing, with insightful direction and invaluable advice. His expertise helped me especially on formulating the research topic and methodology. Professor Xiao Liu provided me with unreserved support since I moved on to the empirical research domain. I benefited a lot from every discussion with her. I also want to express my gratitude to Professor Alan Montgomery and Professor Hui Li. Their expertise and knowledge have made the process of developing this dissertation rewarding and exciting.

I would also like to thank Professor Param Vir Singh. His invaluable advice and support are tremendously helpful for my research and the job market. I learned not only research skills but also precious experience from him. I would also like to thank other professors and friends within and outside Carnegie Mellon University: Nitin Mehta, Dokyun Lee, Yan Huang, Timothy P. Derdenger, Christopher Y. Olivola, Peter Boatwright, Jeff Galak, Liu Liu, Tong (Joy) Lu, Vineet Kumar, and Tong Guo, for taking the time and effort giving advice and discussing research ideas. Their support makes the doctoral study a more rewarding and pleasant journey.

I thank my husband Siyu (Eric) Lu. This dissertation would not have been possible without his endless support, unreserved love, and continued patience. I would also like to express my deepest gratitude to our parents for their love and support. Lastly, I thank my daughter, Yutong

Thea Lu, for the great joy she brought to my life and for the distraction she injected into my work time.

For my dearest family

Contents

Chapter 1

Freemium as an Optimal Strategy for Market Dominant Firms

1.1 Introduction

Over the past years, "freemium" has attracted considerable attention from both practitioners and academics. Many believe that freemium underlies the meteoric rise of companies like Skype and Dropbox, and a horde of startups have jumped on the bandwagon and adopted freemium as their business model of choice. However, successful implementation of the freemium strategy remains challenging. A WSJ report titled "When Freemium Fails" interviewed frustrated entrepreneurs who considered "move(ing) away from freemium" as "the best business decision ... (they) ever made" (Needleman and Loten 2012). An investment manager at First Round Capital summarized the entrepreneurs' frustration as "too many freemium models have too much free and not enough mium". From online gaming to music streaming, leading companies offer widely divergent opinions on whether and when a free option should be offered at all.[1]

Not only is freemium controversial among practitioners, but it also represents a curious case in the eyes of a theoretician. On the surface, freemium resembles a classic case of product-line screening, wherein a firm offers a menu of products at different prices to segment the market. However, as proven by Mussa and Rosen (1978) and more recently by Anderson and Celik (2014), a profit-maximizing firm should always choose inefficient quality but efficient price for its lowest quality product, while doing the reverse for its highest quality product.[2] Said differently, the low-end product's price should be positive and maximize the single-product profit. This theoretical prescription seems to stand in exact opposite to the notion of freemium.

[1]See Zetlin (2013), for example, an interview of Rhapsody's CEO who insisted on its subscription-only model while competitors adopted freemium. In the gaming industry, leading firms such as Blizzard Inc. offer freemium on some of their games but not others.

[2]Throughout the paper, we use the term efficient quality to refer to the quality level that maximizes single product profit (therefore social welfare) under the complete information benchmark.

A number of straightforward reasons come to mind as to why firms find freemium attractive. The illusion of advertising as a "last resort" revenue source is often cited. Saving users the hassle of payment (which can seem high when the price is too small) is another. The power of "free" as a behavioral marketing tool is a third. Although these factors are certainly relevant, we join a nascent literature in marketing and information system that looks at the more nuanced economic reasons behind the freemium phenomenon. In marketing, two pioneering papers by Kamada and Ory (2015) and Lee et al. (2013) have studied the design of freemium to facilitate word-of-mouth and product diffusion. In comparison, we adopt a single-period monopolistic screening framework and study the optimality of freemium when diffusion-related factors are absent. That is, we ask whether and when "perpetual freemium" remains an effective strategy once a product has achieved sufficient recognition and diffusion-related factors have declined in importance.[3] This is especially relevant for firms that have almost reached market saturation. Google Drive and LinkedIn are best examples where word of mouth or diffusion is a non-issue.

More specifically, we ask whether network effects from product usage alone can justify the freemium model, when a firm's sole objective is its single-period product line profit. The notion of network effects speaks to the fact that consumers' valuation of a product varies depending on how many other consumers are using the product or compatible products. Network effects can be generated not only by direct interactions of consumers, but also by indirect behavioral reasons. Direct interaction happens when a free user of Dropbox shares a file with a paid user, when a free player of Farmville trades with a paid user, and when a free user of Spotify shares her playlist. Behaviorally, network effects are created when a consumer values the product more if there are more users of the same product because it allows him/her to socially fit in with their peers, or when a consumer values the paid product more if there are more users of the free product because he/she can derive social prestige from using the high-end product. Intuitively, offering a free version brings more users on board and generates greater network effects. Meanwhile, the Mussa and Rosen (1978)'s insights remain valid and the risk of cannibalization remains high. Ex-ante, it is not clear which is the dominant factor.

By endogenizing a firm's price decision in all relevant subgames, in the baseline model, we build a general framework to study a firm's product line strategy. We pay particular attention to whether and when the firm would *endogenously* choose a zero price for its low-end product. Our first set of results speaks to the conditions under which freemium will *not* hold. We show that, as expected, freemium cannot emerge in a classic screening model without network effects. We are able to prove this with a very general quasi-linear utility function and type distribution, thereby showing the robustness of the Mussa and Rosen (1978) insights.

[3] A number of papers have touched upon this issue (Niculescu and Wu, 2014; Cheng and Tang, 2010), but, none of them have completely endogenized prices and qualities in the product line.

Importantly, even with uniform network effects, freemium remains a dominated strategy. Although network effects lead to stronger incentives to expand the market, they also make the cannibalization effect stronger as more users adopt the low-end ("free") product. When the free product delivers the same network value as the paid product, a price cut on the low-end product will increase the attractiveness of both products by the same margin, thus *tightening* each consumer's incentive compatibility constraint. Thus, more free users does not translate into higher profits from the paid users. Consequently, although network effects give the firm stronger incentives to cover the market, this is done by offering a "conventional product line," wherein the low quality product is priced at a positive level. This result remains valid when we endogenize quality choices for the product line. Put differently, introducing sufficient difference only in "standalone" qualities is not enough to make the freemium strategy viable.

We further show that freemium could indeed become optimal when there is sufficient asymmetry in network effects between the high- and low-end products. In order for freemium to be viable, the firm's product line has to be such that the paid users gain access to larger network effects compared to their non-paying peers. This result somewhat echoes the message in Kumar (2014) that in order to make freemium work, a firm has to offer different sets of features in its free and paid products. But we show that it is the network effects, rather than the "standalone" quality that are the crucial factors.

As an extension, we endogenize the firm's quality decision of products and examine how the optimal quality levels should change with respect to the network effects. To do so, we consider a specific linear utility function and a uniform distribution for consumer type. We show that, in conventional product line design, the optimal quality of the low-end product increases when network effects are higher. Put differently, standalone quality and network effects are complementary to each other. However, in a freemium equilibium, the low-end product's quality decreases when network effects are larger. In other words, low-end product's network effect is a substitute to its standalone quality. This result stems from the fact that in a freemium equilibrium, the entry-level product generates no revenue and its own purpose is to expand a firm's user base. The quality provision should be "just enough" to bring enough users on board.

As a second extension, we fully endogenize the firm's product line decisions – quality, price, as well as network effects – by examining a simpler model where type distribution is discrete. Remarkably, the main insights from the general model remain intact. We compare the qualities of both products provided under freemium and those offered under conventional product line design. Our analysis suggests that, when the firm adopts the freemium strategy, the (non-network) quality gap between the high-end and the low-end products actually shrinks. When adopting freemium strategy, the firm should even provide a low-end product whose quality is above the efficient level. In other words, the firm should offer higher quality and a zero price in order to retain the low-type consumers; this surprising result stands in contrast to the Mussa and Rosen

(1978)'s results of efficient price and inefficient quality. In an optimal product line, quality and network effects are substitutes for the low-end product, but are independent dimensions for the high-end product.

To sum up, our analysis yields a set of managerial recommendations that complement what has been suggested in the previous literature. We show that in the absence of word-of-mouth, "getting more consumers on board" alone cannot justify the freemium strategy. In most cases, market expansion can be more effectively achieved by offering a conventional product line, wherein the low-end product is priced at a positive level to avoid unnecessary cannibalization. In the current framework, perpetual freemium is only optimal under network effects asymmetry. The right freemium strategy should include a free product with lower network benefit than the paid product, but superior standalone functionalities (compared to the efficient level).

The difficulty of establishing freemium arising in equilibrium needs special mention. When consumers derive positive utility for a product, the firm can always charge a positive price and such a deviation is inherently likely to be more profitable. Thus, in general, sustaining freemium in equilibrium is likely to be difficult. Given the extensive and growing nature of freemium, however, demonstrating that such a strategy can arise in equilibrium assumes significant importance.

In the markets where freemium is common, externality benefits are also quite common. Therefore, incorporating externality is a natural and arguably critical element of model formulation to examine product line price and qualities in equilibrium. In spite of a rich model, we get a sharp insight that asymmetric network externality is essential to support freemium. Equally important sharp insight is that freemium will not be sustained when the network effects are the same across all levels of products that differ in quality.

We organize the rest of the paper as follows. In Section 1.2, the related literature and the contribution of the present paper are discussed. Section 1.3 presents the model setup. Section 1.4 presents the analysis and results. Section 1.5 discusses the extension on endogenous quality and network effect decisions, as well as the discrete case. Section 1.6 concludes.

1.2 Literature Review

This paper is related to three streams of literature. First, a number of recent papers have studied various aspects of the freemium strategy. In marketing, Lee et al. (2013) and Kamada and Ory (2015) are among the first studies on the design of freemium. Kamada and Ory (2015) build a micro model of referral behavior and investigate whether a free contract or a referral program is a more efficient means to encourage word of mouth (WOM). The free contract ensures that a receiver would adopt the product even if she turns out to be a low type. When a receiver's adoption generates network effects on the sender, the free contract gives the sender stronger incentives to refer the product in the first place. The main trade-off is between expanded second period demand (due to WOM) and cannibalization. Our model shares some of the features in Kamada and Ory (2015). However, instead of network externalities between first-period senders and second-period receivers, we consider a static model in which network effects exist within and between consumer segments. In our model, an expanded network size leads to the potential to increase the high-end product's price, while a zero price for the low-end product leads to cannibalization. In other words, the focus of this paper is on the optimality of freemium when diffusion dynamics are absent (Mahajan et al., 1990; Chatterjee and Eliashberg, 1990).

In another closely related paper, Lee et al. (2013) develop a structural model to study the design of freemium. Although the paper's focus is empirical, it develops a rich model of consumer behavior that encompasses adoption, upgrade, referral and usage. There are two main differences between Lee et al. (2013) and this paper. First, this paper considers network effects from product usage but does not model diffusion dynamics. Second, Lee et al. (2013) study the design of freemium once the firm has already committed to a zero price for its low-end product. Even though this is a realistic setup in many contexts wherein a firm would commit to freemium for strategic reasons,[4] we are interested primarily in when and why freemium would endogenously emerge to maximize product line profit. Thus, we endogenize the price on the low-end product instead of fixing it to zero.

A number of papers from information systems have studied various aspects of free trial, popular in the software industry (Cheng and Tang, 2010; Niculescu and Wu, 2014; Cheng and Liu, 2012). None of these papers have fully endogenized price in a general model with a general distribution of consumer type. In particular, we allow the low-end product's price to be endogenous.

Therefore, our study is closely related to the rich literature on product line design in both economics and marketing (e.g., Mussa and Rosen 1978; Anderson and Celik 2014; Desai 2001; Desai et al. 2001). We follow the paradigm established in Mussa and Rosen (1978) and consider single-period product line profit as the firm's objective function. The firm chooses how many products to offer and sets a price for each product. As shown more recently in Anderson and Celik (2014), without network effects, the optimal product line problem can be reformulated as a multi-step optimization problem. The firm first chooses the lowest-quality product's price to maximize its revenue, then proceeds to maximize the additional revenue that comes from the second-lowest-quality product. While the standard Mussa and Rosen model does not consider network effects, a number of recent papers have examined the impact of network effects. Jing (2007) examined market segmentation under network externalities and found that the existence of network effects gives the firm stronger incentives to cover the market. The author did not explore the case of freemium, but his main

[4]For example, keeping a product for free saves the need to set up a payment system, potentially lowering a firm's operating cost. Similarly, zero price has been shown to be a particularly powerful marketing tool, and might be preferred to a small but positive price for behavioral reasons.

insights are echoed in this paper.

Finally, in a broad sense, an asymmetric product line is somewhat reminiscent of a two-sided market. In a two-sided market setup, a platform (firm) has incentives to lower the price for one side below the marginal cost as long as this brings value to the other side (e.g., Hagiu 2006; Rochet and Tirole 2006). Although cannibalization is not relevant in a two-sided market context, it is in the context of product line design. We show that it is the coupling of network effects and cannibalization effect that makes the freemium problem unique.

1.3 Model

In analyzing the freemium problem, we intend to make our key insights as general as possible and independent of most specific assumptions on functional form. Thus, we start by presenting a general model in which we make few assumptions on the form of the utility function as well as consumer type distribution. At the same time, we present a running example with linear utility function and uniform distribution of consumer type. This allows us to precisely pin down the conditions for freemium in analytical forms, and we hope that this exercise will strengthen our main intuitions. Next, we present the general model and the running example in turn.

Consider a monopolist who has the option of offering either one or two vertically differentiated products. For notational convenience, we denote the firm's product strategy as Ω. If two products are offered, $\Omega = \{L, H\}$, where L, H stands for the products of relatively low and high quality. If only one product is offered, Ω is a singleton.

There is a unit mass of consumers. They have heterogeneous taste, which is described by the distribution of θ, a density $f(\theta)$ defined on $[0, 1]$. All customers have access to an outside option, the utility of which is u_0. In the case of Dropbox, for example, this outside option denotes the utility a consumer gets from using a traditional form of storage.[5]

For a customer with taste parameter θ, her valuation from consuming product i is $V^i(\theta, D)$, where $i \in \Omega$ and D is the total user base of the firm's product. Let $D = D^{-i} + D^i$ where D^i is the demand for product i, and D^{-i} is the demand for the other type of product if offered. In this framework, $\frac{\partial V^i(\theta,D)}{\partial D}$ captures network benefit for consumers using product i. In other words, we consider a type of "global" network effects where only the total network size affects consumer utility, though the relationship does not have to be linear. The total utility a consumer derives from buying product i is therefore $U^i = V^i(\theta, D) - p^i$, where p^i is the price

[5]In the main text, we consider cases where outside options for all consumers are the same. In Online Appendix B, we provide analysis for the cases where consumers have heterogeneous outside options. In short, heterogeneous outside options do not qualitatively affect our results in the main analysis.

of product i.

We make the following assumptions regarding $V^i(\theta, D)$.

AI. strictly increasing in quality. $\forall \theta, D$,

$$V^H(\theta, D) > V^L(\theta, D)$$

AII. differentiable in D, *and* strictly increasing in D. $\forall \theta, i$,

$$\frac{\partial V^i(\theta, D)}{\partial D} > 0$$

AIII. differentiable in θ, *and* strictly increasing in θ. $\forall D, i$,

$$\frac{\partial V^i(\theta, D)}{\partial \theta} > 0$$

AIV. has increasing differences in θ and quality/network effects. $\forall D$,

$$\frac{\partial\left[V^H(\theta, D) - V^L(\theta, D)\right]}{\partial \theta} > 0$$

The assumption AIV can be alternatively and more restrictively stated as two assumptions that, respectively, speak to the increasing differences conditions regarding consumer type and the standalone product quality/ network effects, i.e., $\frac{\partial\left[V^H(\theta,0)-V^L(\theta,0)\right]}{\partial\theta} > 0$ and $\forall D$, $\frac{\partial\left[V^H(\theta,D)-V^L(\theta,D)\right]}{\partial D\partial\theta} \geq 0$[6]. Here, $V^i(\theta, 0)$ captures the consumer's valuation of product i's quality when consumers do not derive utility from other users (equivalently $D = 0$), and $\frac{\partial V^i(\theta,D)}{\partial D}$ captures the marginal network effect derived from using product i which has user base D. The first inequality implies $V^i(\theta, 0)$ has increasing difference in (θ, i), while the second inequality means the network benefit also has increasing difference in (θ, i). Assumption AIV relaxes the conditions, and requires only $V^i(\theta, D)$ has increasing difference in (θ, i).

The game consists of two stages. First, the firm chooses its menu of products. Next, consumers make purchase decisions conditional upon the firm's menu and belief about all other consumers' decisions. As is typical in a game with network effects, multiple equilibria may exist in the second stage. Specifically, we seek the Nash Equilibrium in the second stage game that is Pareto dominant. Assumption *AIII* guarantees that a consumer of type θ_0 would expect that all other consumers with $\theta > \theta_0$ have a higher evaluation for any product. Therefore, if type θ_0 prefers purchasing product i to the outside option, type $\theta > \theta_0$ would also do so. Similarly, if type θ_0 prefers purchasing the higher-quality product to the lower-quality one, he

[6]These two requirements together are sufficient conditions of assumption *AIV,* please see appendix for proof.

would expect type $\theta > \theta_0$ to do the same. In the appendix, we show that in the baseline model, the Pareto dominant outcome consists of one of two outcomes. In the first scenario, the firm offers the high-quality product only. In the second scenario, the firm offers a product line, with the higher-valuation segment buying the high-quality product and the lower-valuation segment buying the low-quality product.

The assumptions made above are consistent with those in classic papers on product line design, except we also account for network effects. Lemma 1 illustrates how the demand schedules are determined for each product and pricing strategy.

Lemma 1.1. *Assume that both types of products are offered, and the prices are such that both have positive sales. The demand schedule is determined by two marginal consumers at locations θ_L and θ_{HL}, where*

$$V^L(\theta_L, D) - p^L = u_0 \tag{1.1}$$

$$V^H(\theta_{HL}, D) - p^H = V^L(\theta_{HL}, D) - p^L \tag{1.2}$$

with the low-end product serving $[\theta_L, \theta_{HL}]$, and the high-end product serving $[\theta_{HL}, 1]$.

When only the high-end product is offered, the marginal customer type θ_H is determined by

$$V^H(\theta_H, D) - p^H = u_0 \tag{1.3}$$

and only consumers with $[\theta_H, 1]$ are served.

The proof for Lemma 1 is as follows. In choosing between two types of products, a consumer of type θ chooses the low-end product only if $p^H - p^L > V^H(\theta, D) - V^L(\theta, D)$. Given assumption AIV, the larger θ is, the greater the reduction in price is required for a consumer to choose the low-end product. Hence, it is impossible to induce a consumer of type θ_i to purchase a low-end item if the high-end product is purchased by a consumer of type $\theta_j < \theta_i$. From this feature of $V^i(\theta, D)$ and from the fact that the monopolist can make positive profits from serving at least the high-θ consumers, it follows that if both types of products are offered, the monopolist serves $[\theta_L, \theta_{HL}]$ with the low-quality product, and $[\theta_{HL}, 1]$ with the high-quality product, where θ_L denotes the marginal consumer who is indifferent between purchasing the low-end product and the outside option, while θ_{HL} denotes the marginal consumer who is indifferent between purchasing the high-end and low-end product. The firm's profit is thus $\Pi_{HL} = D^H p^H + D^L p^L$. In the case where only the high-end product is offered,[7] the firm serves $[\theta_H, 1]$, where θ_H denotes the marginal customer who is indifferent between purchasing the high-end product and the outside option. In this case, the firm's profit is

[7] Note that it is not an equilibrium wherein the firm sells only the low-end product. Comparatively, offering only the low-end product is dominated and trivial, because the firm can always attract the same user bases with higher price by offering the high-end product. Therefore, selling only the low-end product is trivially dominated.

simply $\Pi_H = D^H p^H$.

A freemium equilibrium is one in which the firm offers both types of products but charges a zero price for the low-end product. It is formally defined as follows:

Definition 1.1. An freemium equilibrium is defined as a product line offering both products H and L, wherein $p^{L*} = 0,\ p^{H*} > 0,\ \Pi^*_{HL} > \Pi^*_H$.

In order to explain our findings more clearly, throughout the paper, we will illustrate our general findings with a running example with valuation function $V^i(\theta, D) = \theta q^i + \alpha^i D$ and uniform distribution of consumer types, i.e., $\theta \sim U[0,1]$. The quality of product i is denoted by q^i, and α^i captures the network benefit one can derive from any other user's usage of product i. This running example satisfies all the assumptions $AI \sim AIV$.

In this running example, the demand for each product and the total demand are given by

$$D^H = \int_{\theta_{HL}}^{1} f(\theta)\, d\theta = 1 - \theta_{HL},$$

$$D^L = \int_{\theta_L}^{\theta_{HL}} f(\theta)\, d\theta = \theta_{HL} - \theta_L,$$

$$D = D^H + D^L = 1 - \theta_L,$$

where the marginal consumers θ_L, θ_{HL} are given by

$$\theta_L q^L + \alpha^L D - p^L = u_0,$$

$$\theta_{HL} q^H + \alpha^H D - p^H = \theta_{HL} q^L + \alpha^L D - p^L.$$

The total profit for the monopolist is thus

$$\Pi_{HL} = p^H D^H + p^L D^L.$$

In the baseline analysis, we therefore endogenize the product set choice as well as the pricing decision. We do not, however, endogenize the quality level nor the level of network effects. In extensions, we endogenize both decisions by considering model formulations that are analytically tractable. According to Definition 1, freemium is different from selecting a price equal to the marginal cost. As demonstrated in extensions (Section 1.5) where quality decision is endogenized, incorporating a positive production cost does not qualitatively affect our results in the main analysis (Section 1.4).

Before we proceed with the analysis, let us briefly discuss our formulation of network effects. We choose a relatively simple formulation wherein each product delivers a different level of network effect (e.g., α_L and α_H in the running example). This is a somewhat standard formulation where network effect is considered as a product attribute. In reality, however, the patterns of network effects can be much more complex. For example, in social games, the network effects are governed by the consumers' obtained utilities when they interact with each other in games. The non-paying users may play at a disadvantage and derive a disutility when interacting with paying users. Though this scenario will indeed imply that the free option generates less overall network effects than the paid option, the exact level of disutility versus utility would depend on the frequency of interaction between the two types of users. In other words, it can be best captured by a case of *local* network effects with four, instead of two, parameters. Throughout the analysis, we choose to present the simplest model where the network effects can be parameterized by two parameters. It should be kept in mind, that our notion of asymmetric network effects can be richer than it seems, capturing cases such as the social game example discussed above. A more detailed analysis that covers the case of quality-dependent and local network effects is in Online Appendix A.

1.4 Analysis

The game has two stages: first, the firm decides whether to offer one or two products, and sets the price for the offered product(s); second, consumers decide on whether to purchase the product that gives the highest utility, or purchase nothing. To explore the conditions under which freemium is an optimal strategy, we discuss three cases in turn:

(1)No network effect $\frac{\partial V^i(\theta,D)}{\partial D} = 0$ for all $i \in \{L, H\}$.

(2)Uniform (symmetric) network effect $\frac{\partial V^i(\theta,D)}{\partial D} = \alpha > 0$ for all $i \in \{L, H\}$, where α is a constant.

(3)Asymmetric network effect $\frac{\partial V^i(\theta,D)}{\partial D}$ differs for $i = L$ and $i = H$.

To prove the optimality of freemium, we consider a necessary condition for freemium: $\frac{\partial \Pi}{\partial p^L}|_{p^L=0} < 0$. In words, the monopoly would like to set a zero price for the low-end product, only when he has no incentives to marginally increase the price when it is already at zero. It turns out that this condition is sufficient to rule out the optimality of freemium in cases (1) and (2), under the general functional form. Let θ_L,θ_H and θ_{HL} be defined by equations $(1), (2), (3)$, and the asterisk $*$ refer to the equilibrium value. Assuming uniform or no network effects, Proposition 1 states the non-existence of freemium under uniform network effects while Corollary 1 illustrates it with the running example.When $\frac{\partial V^i(\theta,D)}{\partial D} = \alpha \geq 0$ for $i \in \{L, H\}$, freemium can never emerge as the equilibrium strategy.

Proposition 1.1. *When* $\frac{\partial V^i(\theta,D)}{\partial D} = \alpha \geq 0$ *for* $i \in \{L, H\}$*, freemium can never emerge as the equilibrium* *strategy.*

For the running example, Proposition 1 can be stated in a more explicite way as in Corollary 1.

Corollary 1.1. *With $\alpha^i = \alpha \geq 0$ for $i \in \{L, H\}$, $V^i(\theta, D) = \theta q^i + \alpha D$ and $\theta \in U[0,1]$, the firm's equilibrium product and pricing strategy can be characterized as follows :*

(a) When $q^L > \alpha > u_0$ and $q^L + u_0 \geq 2\alpha$, the firm offers both products with $p^{L*} = \frac{q^L - u_0}{2}$, $p^{H*} = \frac{q^H - u_0}{2}$, $\theta^*_{HL} = \frac{1}{2}$, $\theta^*_L = \frac{q^L + u_0 - 2\alpha}{2(q^L - \alpha)}$ and $\Pi^*_{HL} = \frac{q^L - u_0}{4}\left(\frac{\alpha - u_0}{q^L - \alpha}\right) + \frac{q^H - u_0}{4}$.

(b) Otherwise, the firm offers only the high-end product, with $p^{H*} = \frac{q^H - u_0}{2}$, $\theta^*_H = \frac{u_0 + q^H - 2\alpha}{2(q^H - \alpha)}$ and $\Pi^*_H = \frac{(q^H - u_0)^2}{4(q^H - \alpha)}$.

We explain the intuition behind Proposition 1 with two sub-cases: one without network effects (i.e., $\frac{\partial V^i(\theta,D)}{\partial D} = 0$ for all $i \in \{L, H\}$), and the other with uniform network effects (i.e., $\frac{\partial V^i(\theta,D)}{\partial D} = \alpha > 0$ for $i \in \{L, H\}$). Although the former is a special case of the latter, separating them helps us build some intuitions. The intuition for the case with zero network effects is consistent with what has been shown in the product line design literature, and has been most succinctly summarized by Anderson and Celik (2014). In a nutshell, setting the low-end product's price to zero generates no revenue *and* puts downward pressure on the high-quality product's price and demand. Thus, the firm can be better off by withdrawing the low-end product altogether. When it does lauch the entry-level product, it always sets a price that is efficient from a single product profit maximization standpoint (see Anderson and Celik 2014 for details).

When network effects are present, how would the firm's optimal product strategy be impacted? More specifically, can network effects lead to an equilibrium wherein the firm pursues the freemium strategy? As discussed in the introduction, increasing network size is a major intuition in favor of the freemium strategy. Although the low-end product generates no revenue and partially cannibalizes the high-end sales, a larger network size brings higher utility to the high-valuation customers, possibly leading to higher price and, therefore, profit. This is akin to the strategy of user subsidization in a two-sided market context. However, a formal analysis reveals that this intuition is not valid when the network benefits for users of both high- and low-end products are positive but identical.

What is the intuition behind? For freemium to be optimal, the optimal price for the low-end product must be zero. A necessary condition for this is $\frac{\partial \Pi_{HL}}{\partial p^L}|_{p^L=0} \leq 0$. In other words, the firm has incentive to decrease the low-end price even it is already at or close to zero. When network effects are symmetric or not too different between the two products, this condition cannot be met. At $p^L = 0$, it is straightforward that $\frac{\partial \Pi_L}{\partial p^L}|_{p^L=0} \geq 0$. Thus, $\frac{\partial \Pi_{HL}}{\partial p^L}|_{p^L=0} \leq 0$ requires that $\frac{\partial \Pi_H}{\partial p^L}|_{p^L=0} \geq 0$. In other words, at an infinitesimal p^L, the profit from the high-end product would increase as the firm decreases its price on the low-end product. A necessary condition for this is that the marginal consumer's incentive compatibility constraint is relaxed as p^L is reduced. This is necessary for either the demand or the price of the high-end product to increase. However, under uniform network effect, this is not possible for the following intuition. As the firm lowers the price for its low-end product, the total user base expands. Due to network effects, each user indeed gains greater surplus from using the high end product. However, because network effects are symmetric, the low-end product also becomes more attractive, by the same margin. In other words, each consumer's incentive compatibility constraint has not been relaxed. The marginal consumer, in fact, now prefers the low-end product more because of a lower price. Thus, the firm cannot charge a higher p^H nor will it have a higher demand from the high-end product. Meanwhile, the firm suffers greater loss from the low valuation segment. Taken together, in the subgame wherein the firm offers two products, decreasing p^L while it is close to zero, always decreases firm profit. Please see the appendix for the formal proof.

In practice, it is certainly rare that different products would deliver exactly the same level of network effects. However, the broad insights remain intact as long as the network benefits of different products are close enough. This corresponds to a wide range of applications where the firm does not or cannot restrict interaction between paid and free users. In most mobile messaging tools, for example, users can send messages to each other regardless of whether they are paying or not. The network aspect of the product is a relatively simple and straightforward feature. It is difficult to restrict the network benefits enjoyed by the non-paying users unless the firm intentionally handicaps the product.

For the case of uniform network effect, Proposition 1 uncovers a fundamental tension between expanding the network size and containing cannibalization. Next, we consider the case in which the firm's high-end and low-end products can deliver different levels of network effects. This is a widely observed practice among firms who successfully pursue the freemium strategy. LinkedIn, for example, gives free users only limited access to view others' profiles, especially contact details. In some of its freemium games, Zynga used to charge "entry tickets" if the users wanted to game with other users. Paid users of Dropbox are able to share more files with others than free users. Proposition 2 states that freemium may indeed emerge as an equilibrium strategy when the high-end and low-end products differ on both the baseline quality as well as the network effects dimension.

Proposition 1.2. *When the following necessary condition is satisfied, freemium can be an equilibrium strategy:*

$$\left[\frac{\partial V^H(\theta_{HL},D)}{\partial D}-\frac{\partial V^L(\theta_{HL},D)}{\partial D}\right]|_{p^{L*}=0,\ p^{H*}} \geq \left[\frac{1+\frac{\int_{\theta_L}^{\theta_{HL}} f(\theta)d\theta}{p^H f(\theta_{HL})}\left(\frac{\partial V^H(\theta_{HL},D)}{\partial\theta_{HL}}-\frac{\partial V^L(\theta_{HL},D)}{\partial\theta_{HL}}\right)}{f\left(\theta_L\right)\frac{\partial\theta_L}{\partial p^L}}\right]|_{p^{L*}=0,\ p^{H*}}$$

Corollary 2 explicitly states the necessary condition for freemium to be optimal for the running example.

Corollary 1.2. *With $\alpha^H \neq \alpha^L$, $V^i(\theta, D) = \theta q^i + \alpha^i D$ for $i \in \{L, H\}$, and $\theta \in U[0,1]$, freemium can be the optimal equilibrium strategy when the following necessary condition is satisfied:*

$$\alpha^H - \alpha^L \geq \left(q^L - \alpha^L\right)\left[1 + \frac{\theta_{HL}^* - \theta_L^*}{p^{H*}}\left(q^H - q^L\right)\right]$$

where $0 < \theta_L^* < \theta_{HL}^* < 1$ *and*

$$\theta_L^* = \frac{u_0 - \alpha^L}{q^L - \alpha^L},$$

$$\theta_{HL}^* = \frac{\left(q^H - q^L\right) q^L + \alpha^H\left(u_0 - q^L\right) - \alpha^L\left(q^H - 2q^L + u_0\right)}{2\left(q^H - q^L\right)\left(q^L - \alpha^L\right)},$$

$$p^{H*} = \frac{q^L\left(\alpha^H - q^L\right) - \left(\alpha^H - \alpha^L\right) u_0}{2\left(q^L - \alpha^L\right)} + \frac{q^H}{2}.$$

From Proposition 2 and Corollary 2, we can see that freemium can be optimal only if the network effect differential $\alpha^H - \alpha^L$ is higher than some positive value.

To solve the profit maximization problem, when two products are offered, first we can get the prices (p^H, p^L) expressed by marginal consumer types (θ_{HL}, θ_L) by Lemma 1. Substituting the price function back to the profit function Π_{HL}, the optimal θ_L^* and θ_{HL}^* are the values that maximize Π_{HL} under the constraint $0 \leq \theta_L < \theta_{HL} < 1$ and $p^L \geq 0$, and then the corresponding p^{L*} and p^{H*} can be obtained. When only the high-end product is offered, we can solve for θ_H^* and thus p^{H*} following similar logic. For the running example, the profit function is concave in θ_k, with $k \in \{H, L, HL\}$; the optimal solutions, thus, can be obtained straightforwardly. The following Corollary 3 gives the optimal product and pricing strategies and therefore describes a sufficient condition for freemium to be optimal.

Corollary 1.3. *With $\alpha^H \neq \alpha^L$, $V^i(\theta, D) = \theta q^i + \alpha^i D$ for $i \in \{L, H\}$, and $\theta \in U[0,1]$, the firm's equilibrium product and pricing strategy can be characterized as follows :*

(a) When $\frac{\partial \Pi_{HL}}{\partial p^L}|_{p^{L*}=0, p^{H*}} \leq 0$, $0 < \frac{u_0 - \alpha^L}{q^L - \alpha^L} < \frac{\left(q^H - q^L\right) q^L - \alpha^L\left(q^H - 2q^L + u_0\right) + \alpha^H\left(u_0 - q^L\right)}{2\left(q^H - q^L\right)\left(q^L - \alpha^L\right)} < 1$, *and* $\Pi_{HL}^* \geq \frac{\left(q^H - u_0\right)^2}{4\left(q^H - \alpha^H\right)}$, *the firm offers two products with*

$$
\begin{aligned}
p^{L*} &= 0, \\
p^{H*} &= \frac{q^L\left(\alpha^H - q^L\right) - \left(\alpha^H - \alpha^L\right)u_0}{2\left(q^L - \alpha^L\right)} + \frac{q^H}{2}, \\
\theta_L^* &= \frac{u_0 - \alpha^L}{q^L - \alpha^L}, \\
\theta_{HL}^* &= \frac{\left(q^H - q^L\right)q^L - \alpha^L\left(q^H - 2q^L + u_0\right) + \alpha^H\left(u_0 - q^L\right)}{2\left(q^H - q^L\right)\left(q^L - \alpha^L\right)}, \\
\Pi_{HL}^* &= \frac{\left[\alpha^L\left(q^H - u_0\right) - q^L\left(q^H - q^L\right) + \alpha^H\left(u_0 - q^L\right)\right]^2}{4\left(q^H - q^L\right)\left(q^L - \alpha^L\right)^2}.
\end{aligned}
$$

(b) When $0 \leq \frac{p^{L*}+u_0-\alpha^L}{q^L-\alpha^L} < \frac{p^{H*}-p^{L*}-\left(\alpha^H-\alpha^L\right)\frac{q^L-p^{L*}-u_0}{q^L-\alpha^L}}{q^H-q^L} < 1$, $0 < p^{L*} < p^{H*}$, *and* $\Pi_{HL}^* \geq \frac{\left(q^H-u_0\right)^2}{4\left(q^H-\alpha^H\right)}$, *the firm offers both products with positive prices and*

$$
\begin{aligned}
\theta_L^* &= \frac{\left(\alpha^H - \alpha^L\right)\left(\alpha^H - \alpha^L - q^H + q^L\right) + 2\left(q^H - q^L\right)\left(\alpha^H + \alpha^L - u_0 - q^L\right)}{\left(\alpha^H - \alpha^L\right)^2 - 4\left(q^H - q^L\right)\left(q^L - \alpha^L\right)} \\
\theta_{HL}^* &= \frac{\left(\alpha^H - \alpha^L\right)^2 + \alpha^L\left(u_0 + 2q^H - 3q^L\right) + \alpha^H\left(q^L - u_0\right) - 2q^L\left(q^H - q^L\right)}{\left(\alpha^H - \alpha^L\right)^2 - 4\left(q^H - q^L\right)\left(q^L - \alpha^L\right)}, \\
\Pi_{HL}^* &= \frac{\left(q^H - q^L\right)\left[\alpha^L\left(q^H - c\right) - \alpha^H\left(q^L - c\right) + \left(2c - q^H\right)q^L - c^2\right]}{\left(\alpha^H - \alpha^L\right) + 4\left(\alpha^L - q^L\right)\left(q^H - q^L\right)}.
\end{aligned}
$$

(c) Otherwise the firm offers only the high-end product with $p^{H*} = \frac{q^H-u_0}{2}$, $\theta_H^* = \frac{q^H+u_0-2\alpha^H}{2(q^H-\alpha^H)}$, $\Pi_H^* = \frac{\left(q^H-u_0\right)^2}{4(q^H-\alpha^H)}$.

Under asymmetric network effects (i.e., $\frac{\partial V^H(\theta_{HL},D)}{\partial D} - \frac{\partial V^L(\theta_{HL},D)}{\partial D}$ has to be greater than a certain positive threshold), the firm can indeed increase total profit by offering the low-end product for free. The key intuition is as follows. As the firm cuts its low-end price to expand demand, both high- and low-end products become more attractive. However, due to the difference in network effects $\frac{\partial V^H(\theta_{HL},D)}{\partial D}$ and $\frac{\partial V^L(\theta_{HL},D)}{\partial D}$, the high-end product becomes relatively more attractive. Put differently, the incentive compatibility constraint for the marginal consumer can be less tight when the network size is larger, due to asymmetric network effects. As such, when $\frac{\partial V^H(\theta_{HL},D)}{\partial D} - \frac{\partial V^L(\theta_{HL},D)}{\partial D}$ is large enough, holding p^H fixed, decreasing p^L may lead to higher demand for the high-end product and therefore higher profit. The premium from the paid users is indeed sufficient to pay for the losses.

We can also see that in the running example, with asymmetric network effects, $q^H - q^L$ needs to be small enough. Intuitively, if $q^H - q^L$ is too large, even with $p^L = 0$, the low-end product quality q^L is not high enough to attract many new customers (i.e., θ_L is too large), and is thus unable to create high enough network benefit. In this case, even a fairly large asymmetry between α^H and α^L cannot induce freemium as the optimal strategy (the necessary condition cannot be satisfied). This condition makes it sufficient for freemium to be adopted as an equilibrium strategy. Please see the appendix for a more detailed discussion.

Figure 1.1: Product line strategy under network effects ($q^H = 0.8$, $q^L = 0.2$, $u_0 = 0.06$)

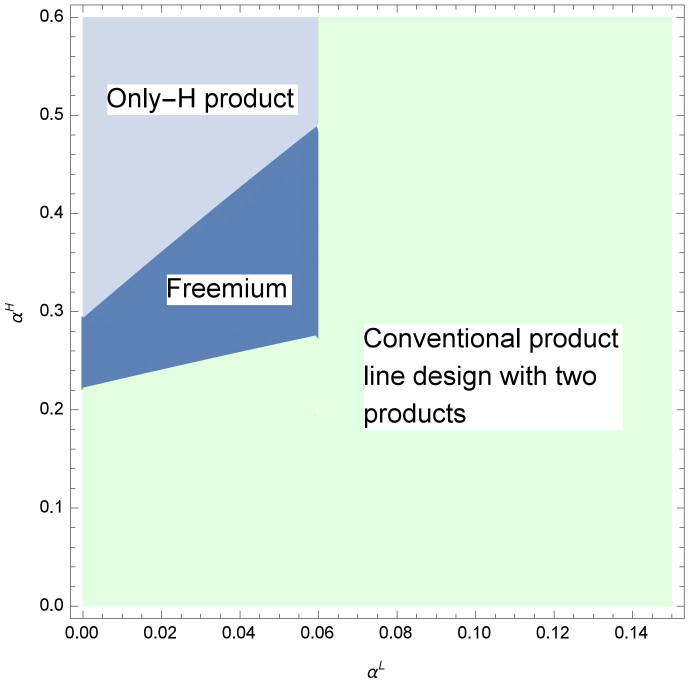

Figure 1.1 illustrates these equilibrium strategies in the parameter space for the running example. Simply put, when the difference between the network effects of low- and high-end products is relatively large, the firm has less pressure from cannibalization and focuses on network effects. This leads to the freemium strategy, wherein the firm forgoes the profit from the low-end customers and considers them a subsidy to the premium users. When the network effects of the low-end product becomes stronger, the firm pursues a conventional product line strategy, where prices for both products are positive. When the high-end product's network effects (i.e. α^H) are very high, the firm would be better off by offering only the high-end product.

It is worth noting that, as shown in Figure 1.1, if we decrease α^L while fixing α^H, the optimal strategy may move from freemium to offering only the high-end product. This may seem puzzling at first glance because a lower α^L leads to a larger asymmetry between the network effects, which further relaxes the cannibalization pressure. However, another effect of lowering α^L is that the firm would find it harder (costlier) to attract consumers to use the low-end product, and therefore, offering only the high-end product becomes more attractive. Mathematically, as shown by the necessary condition in Proposition 2 and Corollary 2, when α^L decreases Δ, the left hand side of the inequality increases Δ, but the right hand side increases more than Δ, making it harder to satisfy the inequality. Similarly, if we increase α^H while fixing α^L, the optimal strategy may also move from freemium to offering only the high-end product. This is because that the profit of offering only the high-end product increases in α^H. Therefore, even though the necessary conditon for freemium to be optimal is easier to satisfy with a larger α^H, the sufficient condition, which requires $\Pi^*_{HL}|_{p^{L*}=0,p^{H*}} > \Pi^*_H$, is harder to satisfy.

Table 1.1 provides a summary of our main insights. We describe the product line choices under each type of equilibrium and provide their existence conditions.

Table 1.1: Summary of product and pricing choices in equilibria

Type of equilibrium	*Features*	*Exists under no n.e.?*	*Exists under uniform n.e.?*	*Exists under asymmetric n.e.?*
(1) *Conventional product line*	*Anderson & Celik* 2014 $p^{H*} > p^{L*} > 0$	✓	✓	✓
(2) *Freemium*	$p^{L*} = 0$, $p^{H*} > 0$	×	×	✓
(3) *Only one product*	$p^* > 0$	✓	✓	✓

It should be noted, that the phenomenon of freemium can be considered as a special case of complementary good pricing. A monopoly who sells two complementary goods has incentives to lower each product's price if this leads to higher sales of the other product. If the complementarity by one product to the other product is stronger, the firm may lower the former product's price below marginal cost in order to profit from the sales of the latter product. Our model essentially builds on this insight in a vertical differentiation scenario. We argue that in a vertical product line, the cannibalization effect co-exists with the complementarity effect. In the case of network effects, the network effects have to be asymmetric so that the complementarity effect outweighs the cannibalization effect when the firm lowers the entry-level product's price. The ecomomic intuition, however, may indeed have broader applications.[8]

1.5 Endogenous quality decisions

The analyses so far speak to the conditions under which freemium is or is not optimal in a product line with given qualities and network effects. In this section, we consider the endogeneous determination of the quality levels and network effects. This is a technically challenging exercise, and we approach it by considering two alternative formulations. Section 1.5.1 considers the running example introduced in Section 1.3, with uniform distribution of consumer type and linear utility function. We allow q^i to be a decision variable and let $C\left(q^i\right)$ denote the firm's marginal production cost of products with quality q^i. This formulation allows us to study

[8] We thank an evaluator for this insightful comment.

the optimal determination of standalone qualities (q^i) when network effects are given. In Section 1.5.2, we fully endogenize both q^i as well as α^i in a discrete segment model, which corresponds to the widely used model of two-type screening.

Remarkably, the key insights in Section 1.4 hold true in all the alternative formulations. Thus, this section could also be considered as a robustness check, while we make more variables endogeneous in progressively simpler models. In both subsections, we relegate the proofs to the appendix and rely on numerical methods to generate the graphical illustrations.

1.5.1 Endogenous quality

With q^H, q^L, p^H, p^L all endogenized, we take a closer look at the equilibrium levels of q^i as a function of network effects. That is, how should the product qualities shift with network effects?

To proceed, we need the following two assumptions in addition to $AI \sim AIV$.

AV. $V^i(\theta, D)$ is concave in q^i, $\forall i \in \{H, L\}$.

AVI. $C\left(q^i\right)$ is convex in q^i, $\forall i \in \{H, L\}$.

With the above two assumptions together with $AI \sim AIV$, we find that the firm should respond to higher α^L by increasing q^L in conventional product line. However, it should respond to higher α^L by reducing quality provision q^L while adopting freemium. Proposition 3 illustrates the firm's choices of quality in the conventional product line and freemium regime, with the running example satisfying $AI \sim AIV$ and marginal cost satisfying *AVI*.

Proposition 1.3. *With valuation function $V^i(\theta, D) = \theta q^i + \alpha^i D$, marginal production cost $C\left(q^i\right) = c \cdot \left(q^i\right)^2$ and $\theta \sim U[0,1]$, when conventional product line is the optimal strategy, the firm offers both products with $q^{H*} = \frac{\theta^*_{HL}}{2c}$, $q^{L*} = \frac{\theta^*_L + \theta^*_{HL} - 1}{2c}$, and q^{L*} increases in α^L; when freemium is the optimal strategy, the firm offers both products with $q^{H*} = \frac{\theta^*_{HLf}}{2c}$, $q^{L*} = \frac{u_0 - \alpha^L\left(1 - \theta^*_{Lf}\right)}{\theta^*_{Lf}}$, and q^{L*} decreases in α^L.*

As explained previously, when a firm pursues the freemium strategy, it suffers a loss on the low-valuation segment and recuperates the lost profit from the high-valuation segment. Because the firm desires the low-valuation consumers for the sole purpose of enlarging its network size, it should always supply the least level of quality that is sufficient to induce purchase. As α^L increases, the minimal required quality level decreases accordingly, leading to lower quality provision. Broadly speaking, quality and network effects are substitutes for the low-end product. The firm always seeks the least costly way to attract the low-end customers, and as one dimension gets higher, it decreases its investment in the other dimension.

1.5.2 Endogenous quality and network effect

In the baseline model, we examined a general utility functional form as well as a general consumer distribution. In Section 1.5.1, the product quality is endogenized. Next, we endogenize both quality and network effect decisions. In reality, the network effect—as a product attribute—may also be the firm's endogenous decision. For example, in social games, the game designer (the firm) can endogenously decide how much network effect the paid users and free users can get, by designing the frequency of interaction between different types of players. In the case of Dropbox, the firm can make sharing more or less convenient so that different products deliver different network effect.

To keep the analysis tractable, we consider a discrete distribution of consumers on the demand side. Namely, there are two segments of consumers, with high and low valuation for the product quality as well as network benefit. Each consumer is characterized by a taste parameter $\theta \in \{\theta_H, \theta_L\}$, where $\theta_H > \theta_L$. A fraction λ of consumers belong to type θ_H, who have higher valuation for the firm's products. A fraction $1-\lambda$ of consumers belong to type θ_L. There may be some debate on the formulation of the running example regarding whether the taste parameter (i.e., θ) affects only the valuation of the standalone quality, or the valuation of both quality and network benefit. In short, both formulations satisfy $AI \sim AIV$, therefore, the results derived for the general model in Section 1.4 hold for both specifications. To demonstrate our results hold for both cases, we therefore offer a formulation in this extension different from that used in the previous running example. For a customer with taste parameter θ, her valuation from consuming product i is:

$$V^i(\theta, \alpha^i, D) = \theta(q^i + \alpha^i D),$$

where $D \leq 1$ is the total number of users who buy from the firm's product line, namely $D = \sum_{i \in \{H,L\}} D^i$. As such, we assume that each user derives network effects from all other users in the firm's network. The total magnitude of network effects depends on the network size and product design. The firm sets price $p^i > 0$ for each product i. To guarantee the existence of interior solutions, we assume that the marginal cost of serving a consumer is increasing in both q^i and α^i quadratically, i.e., $C\left(q^i, \alpha^i\right) = c\left(q^i\right)^2 + s\left(\alpha^i\right)^2$. The firm's product line profit is thus:

$$\Pi = \sum_{i \in \{H,L\}} D^i \left[p^i - c\left(q^i\right)^2 - s\left(\alpha^i\right)^2\right].$$

We can see that all of $AI \sim AVI$ are satisfied. In this discrete case, we also assume that the high-type consumers have positive valuation for the low-end product at price zero.[9] Proposition 4 spells out the optimal

[9] In appendix B, we explain specifically for the discrete case where this assumption does not hold, which means that the highest-type consumers will not find the low-end product worth trying even if it is offered at zero price. This is implausible if

product line design under uniform network effect.

Proposition 1.4. *When $\alpha^H = \alpha^L = \alpha$, the equilibrium product-line strategy can be characterized as follows:*

$$\text{When } (I)\begin{cases} \frac{\theta_L^2-\lambda^3\theta_H^2}{4s} - (1-\lambda)u_0 > 0 \\ \lambda\theta_H\theta_L > \theta_L^2 \geq 2su \end{cases} \text{ or } (II)\begin{cases} \frac{(\lambda\theta_H-\theta_L)^2}{4c(1-\lambda)} + \frac{2\theta_L-\lambda^3\theta_H^2}{4s} - (1-\lambda)u_0 > 0 \\ s\left[\theta_H\theta_L\lambda + 2cu_0(1-\lambda)\right] \leq \theta_L^2(c - c\lambda + s) \\ \lambda \leq \frac{\theta_L}{\theta_H} \end{cases} \text{ the firm offers both}$$

products with price $p^{H*} > p^{L*} > 0$.[10] *The corresponding optimal qualities, network intensity and firm profit are:*

$$q^{L*} = \begin{cases} 0 & ,\textit{when } (I) \textit{ holds} \\ \frac{\theta_L-\theta_H\lambda}{2c(1-\lambda)} & ,\textit{when } (II) \textit{ holds} \end{cases}$$

$$q^{H*} = \frac{\theta_H}{2c},\ \alpha^* = \frac{\theta_L}{2s}$$

$$\Pi_{HL}^* = \begin{cases} \frac{\lambda\theta_H^2}{4c} + \frac{\theta_L^2}{4s} - u_0 & ,\textit{when } (I) \textit{ holds} \\ \frac{\theta_L^2+\lambda\theta_H^2-2\lambda\theta_H\theta_L}{4c(1-\lambda)} + \frac{\theta_L^2}{4s} - u_0 & ,\textit{when } (II) \textit{ holds} \end{cases}$$

Otherwise, the firm offers only a high-end product. The price, quality, and equilibrium profit are:

$$p^* = \frac{\theta_H^2}{2c} + \frac{\theta_H^2\lambda^2}{2s} - u_0,$$
$$q^* = \frac{\theta_H}{2c},\ \alpha^* = \frac{\lambda\theta_H}{2s},$$
$$\Pi_H^* = \lambda\left(\frac{\theta_H^2}{4c} + \frac{\lambda^2\theta_H^2}{4s} - u_0\right).$$

The following Figure 1.2 shows regions for optimal product line strategies under zero network effects as well as uniform network effects. The figure 1.2(a) corresponds to the case where $\alpha = 0$, i.e., consumers derive no network benefit from using the product. The figure 1.2(b) depicts the regions for optimal strategies for the case with $\alpha > 0$, namely, the network effects are positive and symmetric.

From the above figures, we can see that freemium is never an optimal strategy when network effects are zero or symmetric, consistent with the results of our baseline model.

Next, we consider the case in which the firm can design its high-end and low-end products to deliver different levels of network effects. Proposition 5 states the optimal product line design when the network effects can be designed as asymmetric.

Proposition 1.5. *When the firm can set α^H and α^L at different levels, the equilibrium product-line strategy can be characterized as follows:*

not impossible according to those successful freemium products offered in the market.

[10] The exact expressions of optimal prices are provided in the appendix.

Figure 1.2: Regions for optimal strategies ($c = 0.5, s = 0.1, \theta_H = 5, \theta_L = 2$)

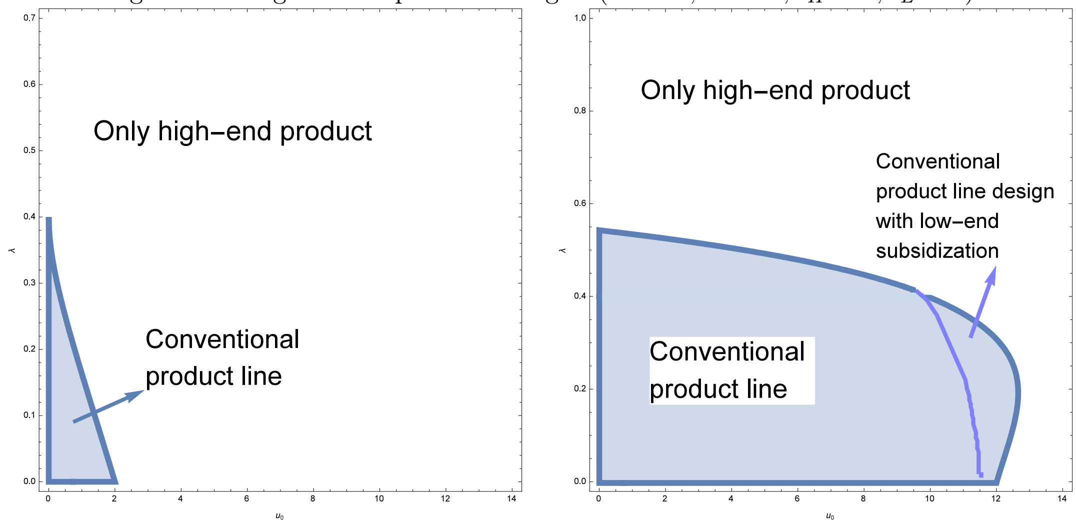

(a) Case without network effect (b) Case with uniform network effect

When $\frac{\theta_H^2}{4s}(\lambda - \lambda^3) + \lambda u_0(1 - \frac{\theta_H}{\theta_L}) - \frac{csu_0^2(1-\lambda)}{\theta_L^2(c+s)} \geq 0$ *and* $\frac{2csu_0(1-\lambda)}{\theta_L{}^2(c+s)} + \frac{\lambda\theta_H}{\theta_L} - 1 > 0$*, the firm adopts freemium strategy with:*

$$p^{L*} = 0,\ p^{H*} = \frac{\theta_H^2}{2}\left(\frac{1}{s} + \frac{1}{c}\right) - \frac{u_0\theta_H}{\theta_L}$$

The corresponding optimal qualities, network effects and profit are:

$$q^{L*} = \frac{su_0}{\theta_L\,(c+s)},\ q^{H*} = \frac{\theta_H}{2c}$$
$$\alpha^{L*} = \frac{cu_0}{\theta_L\,(c+s)},\ \alpha^{H*} = \frac{\theta_H}{2s}$$
$$\Pi_F^* = \frac{\theta_H^2\lambda}{4}\left(\frac{1}{s} + \frac{1}{c}\right) - \frac{cs\,(1-\lambda)\,u_0^2}{\theta_L^2\,(c+s)} - \frac{\lambda u_0\theta_H}{\theta_L}$$

When $\frac{2csu_0(1-\lambda)}{\theta_L{}^2(c+s)} + \frac{\lambda\theta_H}{\theta_L} - 1 \leq 0$ *and* $\frac{(\theta_L-\theta_H\lambda)^2}{4c(1-\lambda)}\left(\frac{1}{s} + \frac{1}{c}\right) + \frac{\theta_H^2}{4s}\left(\lambda - \lambda^3\right) - u_0\,(1-\lambda) > 0$*, the firm launches two products with* $p^{H*} > p^{L*} > 0$*. The corresponding qualities, network effects and profit are:*

$$q^{L*} = \frac{\theta_L - \theta_H\lambda}{2c(1-\lambda)},\ q^{H*} = \frac{\theta_H}{2c},$$
$$\alpha^{L*} = \frac{\theta_L - \theta_H\lambda}{2s(1-\lambda)},\ \alpha^{H*} = \frac{\theta_H}{2s},$$
$$\mathit{\Pi}_{HL}^* = \frac{\theta_L^2 + \theta_H^2\lambda - 2\theta_L\theta_H\lambda}{4c(1-\lambda)} + \frac{\theta_H^2\lambda}{4s} + \frac{(\theta_L - \theta_H\lambda)^2}{4s\,(1-\lambda)} - u_0.$$

Otherwise, the firm provides only the high-end product, with price, quality, network effect and profit as:

$$
\begin{aligned}
p^* &= \frac{\theta_H^2}{2c} + \frac{\lambda^2\theta_H^2}{2s} - u_0, \\
q^* &= \frac{\theta_H}{2c},\ \alpha^* = \frac{\lambda\theta_H}{2s}, \\
\Pi_H^* &= \lambda\left(\frac{\theta_H^2}{4c} + \frac{\lambda^2\theta_H^2}{4s} - u_0\right).
\end{aligned}
$$

Figure 1.3: Strategies with asymmetric network effects ($c = 0.5, s = 0.1, \theta_H = 5, \theta_L = 2$)

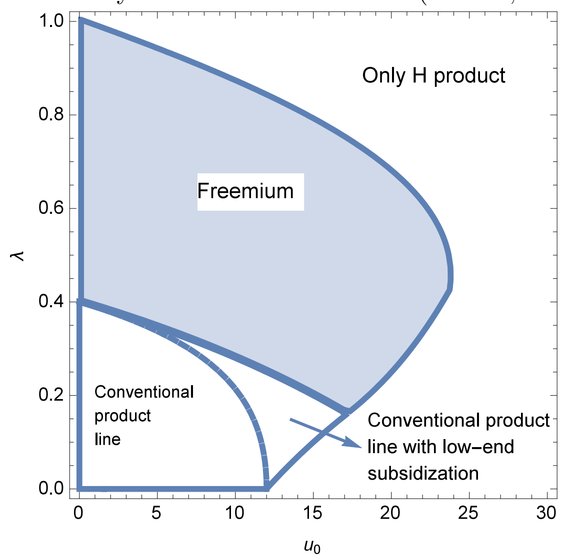

From Figure 1.3, we can see that under asymmetric network effects, the firm can indeed increase product-line profit by offering the low-end product for free. However, it should be noted that even when asymmetric network effects are present, freemium is not always an optimal strategy. When both λ and u_0 are relatively low, the firm segments the market via a conventional product line. In this case, p^L once again corresponds to the efficient price.

Given the above results, we compare the optimal quality decision under freemium and conventional product line design. We use the term "efficient quality", denoted by q^{i0*} with $i \in \{H, L\}$, to refer to the quality level that maximizes single product profit (therefore social welfare) under the complete information benchmark.

Corollary 1.4. *Across all equilibria, the quality of the high-end product is always set at the efficient level, i.e.* $q^{H*} = q^{H0*}$;

When both segments are served with positive prices, the quality of the low-end product is always below the efficient level, i.e. $q^{L*} < q^{L0*}$;

When the freemium strategy is adopted, the quality of the low-end product can be lower, equal to, or even *greater than the efficient level, i.e.,* $q^{L*} = \frac{su_0}{\theta_L(c+s)} > q^{L0*}$ *can hold;*

The quality difference $\Delta q = q^{H*} - q^{L*}$ *is smaller under the freemium strategy than that under conventional product line strategy with uniform network effects.*[11]

For freemium to be optimal, previous studies have recognized the importance of offering a balanced set of features in a firm's free product (Lee et al., 2013; Kumar, 2014). Kumar (2014) stated that for freemium to work, the free offering has to be "compelling enough" to attract new users, but it cannot be "too rich" such that people stick to the free product. This insight is confirmed by our analysis. Whenever a freemium strategy is pursued, the firm has to strike a balance between getting consumers on board and minimizing cannibalization.

At the same time, this insight alone is not enough for the design of an optimal product line. Corollary 4 speaks to the importance of distinguishing between features that contribute to a product's standalone quality and features that contribute to its network effects. It states that when a firm pursues the freemium strategy, it should choose a standalone quality level for its low-end product that is higher than what it would choose in a conventional product line. In other words, a firm pursuing the freemium strategy should offer more features at a lower price (i.e., zero)! This reduces the quality differentiation within the product line. To compensate, the firm should design the products such that they offer different levels of network effects. In other words, it is not just the number of free features that determines freemium's viability, but also which features are included in the free version. The prescription of higher standalone quality may seem counterintuitive at first, but a firm should realize that it is precisely a high standalone quality that allows the firm to choose lower network effects for the low-end product. Lower network effects are the key to minimizing cannibalization.

Figure 1.4 provides a simple illustration of this idea. It plots the equilibrium level of q^{L*} against the value of u_0. As u_0 increases, the equilibrium shifts from a conventional product line to the freemium equilibrium, and q^{L*} increases accordingly.

1.6 Concluding Remarks

This paper studies a monopolist's product and pricing strategy under network effects. We are particularly interested in the optimality of the freemium strategy. We seek to answer two questions. First, what are the necessary conditions for the optimality of freemium? Our results point to the asymmetry in network effects as the determining factor. Second, what are the principles that should guide the design of freemium? Our results add to the previous literature and speak to the distinction between a product's "baseline" quality and

[11] Here we compare the quality difference, assuming that the relevant parameter values are the same.

Figure 1.4: q^{L*} as a function of u_0

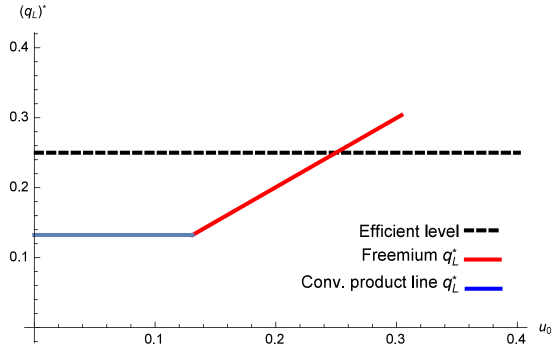

the network effects its users receive. Compared to a conventional product line, in a freemium equilibrium, the firm should offer relatively higher baseline quality but low network effects on its low-end product.

Of course, as we have reviewed in Section 1.2, our analysis provides only one of many explanations of the optimality of freemium. Kamada and Ory (2015), for example, focused on the alternative explanation where freemium motivates word-of-mouth during the diffusion process. Other papers we reviewed in Section 1.2 considered the importance of consumer learning and free trial. In addition, competition may also play a key role in driving the price of the low-end product to the marginal cost. These explanations are clearly not mutually exclusive. Which explanation serves as the most likely explanation to the observed freemium depends on the context that is being considered. Our framework applies best to a scenario where the product is beyond the diffusion stage, and competition is not intense. When competition is relevant, the same mechanism may still be at work, but competition itself creates a downward pressure on the prices and provides the focal firm strong incentives to pull out the low-end product altogether. The exact effect is an interesting topic for future research.

Freemium has become an immensely popular business model among start-ups, especially in the Internet sector. Without doubt, providing a product for free is an effective way to expand a firm's user base. As many entrepreneurs have rightly believed, expanding a firm's user base is of ultimate importance for industries with network effects. However, our analysis points out that a firm should exercise caution when it is tempted to jump on the freemium bandwagon in order to "get users on board." For freemium to work, a firm has to understand the subtleties of network effects in its market. When network effects are not strong enough or are uniform across segments, offering an entry level product for free does no good to the product line profit. Freemium expands a firm's market share but severely limits its ability to create enough margin from its paid users. This scenario may sound familiar to many firms who are frustrated by the disappointing number of paying users under their freemium strategy. Instead, these firms should heed the wisdom of Mussa and Rosen (1978), and segment the market via a conventional product line. A conventional product line consists of a low quality product that is sold at an efficient price, instead of an entry level product sold at zero price.

Our results in the extensions also provide guidelines to firms who should certainly adopt the freemium strategy. As previous studies have pointed out, a firm should provide an intermediate number of features in its free product. Moreover, it is not just the number of features that matters, but also which features the firm should provide. Our analysis prescribes that a firm should in fact be quite generous with features that enhance a product's "baseline quality" – that is, the value of a product when it is used alone. However, the firm should deliberately limit the features that bring network benefits to users. In a well designed freemium menu, "paying for upgrade" should in fact be "paying for network effects."

Our study focuses on a scenario wherein product line profit is the main driver of firm strategy. In doing so, we leave out many behavior factors that are nonetheless relevant to the freemium strategy. For example, offering a product for free can induce greater word-of-mouth and speed up its adoption (e.g., Kamada and Ory 2015). It would be interesting to combine the two perspectives and investigate the dynamics of the freemium strategy. This exercise may lead to a "taxonomy" of the freemium strategy and elucidate the possible motivations behind offering a free product. Second, it is interesting to extend the current model and exam the possibilities of advertising income. Advertising income should give the firm stronger incentives to provide the free product, and a firm should strike a balance between lower fees and higher ad revenue. Third, competition is a relevant factor in many high tech markets. Even though product-line competition brings considerable complexities to the model, it remains a meaningful direction for future research. Finally, it is of some interest to generalize the model into a case of user subsidization, where negative price is possible. When the firm is able to subsidize the users, the price for the entry-level product is not constrained to be non-negative. Although it is relatively easy to persuade adoption with subsidies, it is much harder to induce actual usage. It is of managerial interests to explore strategies that the firm could follow when network effects stem mostly from actual usage.

Reference

Anderson, S. and Celik, L. (2014). Product line design. *Working Paper*.

Banerji, A. and Dutta, B. (2009). Local network externalities and market segmentation. *International Journal of Industrial Organization*, 27(5):605–614.

Cheng, H. K. and Tang, Q. C. (2010). Free trial or no free trial: optimal software product design with network effects. *European Journal of Operational Research*, 205(2):437–447.

Church, J. and Gandal, N. (1992). Network effects, software provision, and standardization. *The Journal of Industrial Economics*, 40(1):85–103.

Desai, P., Kekre, S., Radhakrishnan, S., and Srinivasan, K. (2001). Product differentiation and commonality in design: balancing revenue and cost drivers. *Management Science*, 47(1):37–51.

Desai, P. S. (2001). Quality segmentation in spatial markets: when does cannibalization affect product line design? *Management Science*, 20(3):265–283.

Dey, D., Lahiri, A., and Liu, D. (2013). Consumer Learning and time-locked trials of software products. *Journal of Management Information Systems*, 30(2):239–268.

Farrell, J. and Klemperer, P. (2007). Coordination and lock-in: competition with switching costs and network effects. *Handbook of Industrial Organization*, 3(Chapter 31):1970–1994.

Fjeldstad, Ø., Riis, C., and Moen, E. R. (2010). Competition with local network externalities. *Working Paper*.

Hagiu, A. (2006). Pricing and commitment by two-sided platforms. *The RAND Journal of Economics*, 37(3):720–737.

Jackson, M. O. and Wolinsky, A. (1996). A strategic model of social and economic networks. *Journal of Economic Theory*, 71(0108):44–74.

Jing, B. (2007). Network externalities and market segmentation in a monopoly. *Economics Letters*, 95(1):7–13.

Jing, B. and Zhang, Z. J. (2011). Product line competition and price promotions. *Quantitative Marketing and Economics*, 9(3):275–299.

Jones, R. (2015). Freemium debate heats up. *Music Week*.

Kamada, Y. and Ory, A. (2015). Encouraging word of mouth : Free contracts , referral programs , or both ? *Working Paper*.

Katz, M. and Shapiro, C. (1985). Network externalities, competition, and compatibility. *American Economic Review*, 75:424–440.

Katz, M. L. and Shapiro, C. (1986). Technology adoption in the presence of network externalities. *Journal of Political Economy*, 94:822–841.

Kumar, V. (2014). Making "freemium" work. *Harvard Business Review*, (May):701–703.

Lee, C., Kumar, V., and Gupta, S. (2013). Designing freemium: A model of consumer usage, upgrade, and referral dynamics. *Working Paper*.

Maskin, E. and Riley, J. (1984). Monopoly with incomplete information. *The RAND Journal of Economics*, 15(2):171–196.

Moorthy, K. (1988). Product and price competition in a duopoly. *Marketing Science*, 7(2):141–168.

Moorthy, K. and Png, I. (1992). Market segmentation, cannibalization, and the timing of product introductions. *Management Science*, 38(3):345–359.

Mussa, M. and Rosen, S. (1978). Monopoly and product quality. *Journal of Economic Theory*, 18:301–317.

Needleman, S. and Loten, A. (2012). When freemium fails. *The Wall Street Journal*.

Niculescu, M. and Wu, D. (2014). Economics of free under perpetual licensing: Implications for the software industry. *Information Systems Research*, 25(1):173–199.

Orhun, A. (2009). Optimal product line design when consumers exhibit choice set-dependent preferences. *Marketing Science*, 28(5):868–886.

Riggins, F. (2003). Market segmentation and information development costs in a two-tiered fee-based and sponsorship-based web site. *Journal of Management Information Systems*, 19(3):69–86.

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

Shampanier, K., Mazar, N., and Ariely, D. (2007). Zero as a special price: The true value of free products. *Marketing science*, 26(6):742–757.

Varian, H. R. and Shapiro, C. (1998). Versioning: The smart way to sell information. *Harvard Business Review*, 107(6):107.

Wagner, T. M., Benlian, A., and Hess, T. (2014). Converting freemium customers from free to premium–the role of the perceived premium fit in the case of music as a service. *Electronic Markets*, 24(4):259–268.

Wei, X. D. and Nault, B. R. (2013). Experience information goods: "Version-to-upgrade". *Decision Support Systems*, 56:494–501.

Zetlin, M. (2013). Rhapsody CEO: Freemium doesn't work. *www.inc.com*.

Zhang, K. and Sarvary, M. (2015). Differentiation with user-generated content. *Management Science*, 61(4):898–914.

Appendix A: Proofs.

Proof. ***Footnote 5.***

Below we prove $\forall D$, $\frac{\partial\left[V^H(\theta,D)-V^L(\theta,D)\right]}{\partial\theta}>0$ given the following two conditions:

$$\forall\theta, \frac{\partial\left[V^H(\theta,0)-V^L(\theta,0)\right]}{\partial\theta}>0(*) \text{ and } \forall\theta, D, \frac{\partial\left[\frac{\partial V^H(\theta,D)}{\partial D}-\frac{\partial V^L(\theta,D)}{\partial D}\right]}{\partial\theta}\geq 0(**).$$

We can write

$$\begin{aligned} V^H(\theta,D)-V^L(\theta,D) &= \left[V^H(\theta,0)+\int_0^D\frac{\partial V^H(\theta,t)}{\partial t}dt\right]-\left[V^L(\theta,0)+\int_0^D\frac{\partial V^L(\theta,t)}{\partial t}dt\right] \\ &= V^H(\theta,0)-V^L(\theta,0)+\int_0^D\frac{\partial\left[V^H(\theta,t)-V^L(\theta,t)\right]}{\partial t}dt \end{aligned}$$

then given $(*)$ and $(**)$ we have

$$\begin{aligned} \frac{\partial\left[V^H(\theta,D)-V^L(\theta,D)\right]}{\partial\theta} &= \frac{\partial\left[V^H(\theta,0)-V^L(\theta,0)+\int_0^D\frac{\partial\left[V^H(\theta,t)-V^L(\theta,t)\right]}{\partial t}dt\right]}{\partial\theta} \\ &= \frac{\partial\left[V^H(\theta,0)-V^L(\theta,0)\right]}{\partial\theta}+\int_0^D\frac{\partial\left[\frac{\partial V^H(\theta,D)}{\partial D}-\frac{\partial V^L(\theta,D)}{\partial D}\right]}{\partial\theta}dt \\ &> 0 \end{aligned}$$

□

Proof. ***Lemma 1.1.***

We used proof by contradiction for Lemma 1. Please refer to the main text for the complete proof. □

Proof. ***Proposition 1.1.***

To prove Proposition 1, we show that when $\frac{\partial V^H(\theta,D)}{\partial D}=\frac{\partial V^L(\theta,D)}{\partial D}\geq 0$, the necessary condition for freemium to be optimal (i.e., $\frac{\partial\Pi_{HL}}{\partial p^L}|_{p^L=0}\leq 0$) is violated.

Let us now derive the general expression of Π as a piecewise function of p^H and p^L. Consider θ_L,θ_H and θ_{HL} which are determined by

$$V^L(\theta_L, D) - p^L = u_0$$

$$V^H(\theta_{HL}, D) - p^H = V^L(\theta_{HL}, D) - p^L$$

$$V^H(\theta_H, D) - p^H = u_0$$

Clearly, θ_L,θ_H and θ_{HL} are implicit functions of p^H and p^L. As a consequence, the profit function Π is a function of p^H and p^L:

$$\Pi = \begin{cases} p^H & \text{, when } p^H \text{and } p^L \text{are s.t. } \theta_L \leq 0 \ \& \ \theta_{HL} \leq 0. \\ & \text{In this case } D^L = 0, D^H = 1. \\ p^L & \text{, when } p^H \text{and } p^L \text{are s.t. } \theta_L \leq 0 \ \& \ \theta_{HL} \geq 1. \\ & \text{In this case } D^L = 1, D^H = 0. \\ 0 & \text{, when } p^H \text{and } p^L \text{are s.t. } \theta_L \geq 1 \ \& \ \theta_H \geq 1. \\ & \text{In this case } D^L = D^H = 0. \\ p^H D^H & \text{, when } p^H \text{and } p^L \text{are s.t. } \theta_L \geq 1 \ \& \ \theta_H \in (0,1) \text{, } or \ 0 < \theta_H < \theta_L \leq 1. \\ & \text{In this case } \ D^L = 0, D^H = \int_{\theta_H}^1 f(\theta)\, d\theta. \\ p^L D^L & \text{, when } p^H \text{and } p^L \text{are s.t. } \theta_L \in (0,1) \ \& \ \theta_{HL} \geq 1. \\ & \text{In this case } D^L = \int_{\theta_L}^1 f(\theta)\, d\theta, D^H = 0. \\ p^L D^L + p^H D^H & \text{, when } p^H \text{and } p^L \text{are s.t. } \theta_L \leq 0 \ \& \ \theta_{HL} \in (0,1) \text{, } or \ 0 < \theta_L < \theta_{HL} \leq 1. \\ & \text{In this case } D^L = \int_{\theta_L \cdot I(\theta_L \geq 0)}^{\theta_{HL}} f(\theta)\, d\theta, D^H = \int_{\theta_{HL}}^1 f(\theta)\, d\theta. \end{cases}$$

$$= \begin{cases} p^H & \text{, when } p^H \text{and } p^L \text{are s.t. } \theta_L \leq 0 \ \& \ \theta_{HL} \leq 0. \\ p^L & \text{, when } p^H \text{and } p^L \text{are s.t. } \theta_L \leq 0 \ \& \ \theta_{HL} \geq 1. \\ 0 & \text{, when } p^H \text{and } p^L \text{are s.t. } \theta_L \geq 1 \ \& \ \theta_H \geq 1. \\ p^H \int_{\theta_H}^1 f(\theta)\, d\theta & \text{, when } p^H \text{and } p^L \text{are s.t. } \theta_L \geq 1 \ \& \ \theta_H \in (0,1)\text{,} \\ & \ or \ 0 < \theta_H < \theta_L \leq 1. \\ p^L \int_{\theta_L}^1 f(\theta)\, d\theta & \text{, when } p^H \text{and } p^L \text{are s.t. } \theta_L \in (0,1) \ \& \ \theta_{HL} \geq 1. \\ p^L \int_{\theta_L \cdot I(\theta_L \geq 0)}^{\theta_{HL}} f(\theta)\, d\theta + p^H \int_{\theta_{HL}}^1 f(\theta)\, d\theta & \text{, when } p^H \text{and } p^L \text{are s.t. } \theta_L \leq 0 \ \& \ \theta_{HL} \in (0,1)\text{,} \\ & \ or \ 0 < \theta_L < \theta_{HL} \leq 1. \end{cases}$$

First notice that, in equilibrium, θ_L^* cannot be smaller than 0, because the firm can always increase profit by increasing both p^H and p^L to make the type $\theta = 0$ indifferent between purchasing and not purchasing. Thus, we need to prove $\Pi_{HL} = p^L \int_{\theta_L}^{\theta_{HL}} f(\theta)\, d\theta + p^H \int_{\theta_{HL}}^1 f(\theta)\, d\theta$ violates the necessary condition $\frac{\partial \Pi_{HL}}{\partial p^L}|_{p^L=0} \leq 0$, for any p^H as long as the demand schedule is $0 \leq \theta_L < \theta_{HL} \leq 1$.

Taking partial derivative of Π_{HL} w.r.t. p^L, we have

$$\frac{\partial \Pi_{HL}}{\partial p^L} = \int_{\theta_L}^{\theta_{HL}} f(\theta)\, d\theta + p^L \left[f(\theta_{HL}) - f(\theta_L)\right] \left(\frac{\partial \theta_{HL}}{\partial p^L} - \frac{\partial \theta_L}{\partial p^L}\right) - p^H f(\theta_{HL}) \frac{\partial \theta_{HL}}{\partial p^L}$$

The sign of $\frac{\partial \Pi_{HL}}{\partial p^L}$ depend on the sign of $\frac{\partial \theta_L}{\partial p^L}$ and $\frac{\partial \theta_{HL}}{\partial p^L}$. Next we determine the signs of $\frac{\partial \theta_L}{\partial p^L}$ and $\frac{\partial \theta_{HL}}{\partial p^L}$.

Recall from equation (1) , θ_L is implicitly given by $V^L(\theta_L, D) = p^L + u_0$. From assumption *AII* and *AIII*, we have $\frac{\partial V^i(\theta,D)}{\partial \theta} > 0$, $\frac{\partial V^i(\theta,D)}{\partial D} > 0$, and we also have D decreasing in p^L. Therefore, when p^L increases, we must have a higher θ_L to maintain the equality $V^L(\theta_L, D) = p^L + u_0$. In other words, we can get

$$\frac{\partial \theta_L}{\partial p^L} > 0 \tag{1.4}$$

Recall from equation (2) , θ_{HL} is implicitly given by $V^H(\theta_{HL}, D) - V^L(\theta_{HL}, D) = p^H - p^L$. When $\frac{\partial V^i(\theta,D)}{\partial D} = 0$ or $\frac{\partial V^H(\theta,D)}{\partial D} = \frac{\partial V^L(\theta,D)}{\partial D}$, we have $V^H(\theta_{HL}, D) - V^L(\theta_{HL}, D) = V^H(\theta_{HL}, 0) - V^L(\theta_{HL}, 0)$. Namely, the demand change does not affect the valuation differential. So we have $V^H(\theta_{HL}, 0) - V^L(\theta_{HL}, 0) = p^H - p^L$. From assumption *AIV*, we have $\frac{\partial \left[V^H(\theta,0) - V^L(\theta,0)\right]}{\partial \theta} > 0$. Therefore when p^L decreases, we must have a higher θ_{HL} to maintain the equality $V^H(\theta_{HL}, 0) - V^L(\theta_{HL}, 0) = p^H - p^L$. In other words, we can get

$$\frac{\partial \theta_{HL}}{\partial p^L} < 0 \tag{1.5}$$

Because $\theta_{HL} > \theta_L$ and $f(\theta_{HL}) \geq 0$, with (5) we have

$$\begin{aligned}\frac{\partial \Pi_{HL}}{\partial p^L}|_{p^L=0} &= \int_{\theta_L}^{\theta_{HL}} f(\theta)\, d\theta - p^H f(\theta_{HL}) \frac{\partial \theta_{HL}}{\partial p^L}|_{p^L=0} \\ &> 0\end{aligned}$$

□

Proof. ***Corollary 1.1.***

If the firm offers both products, according to equations (1) and (2), we have

$$\theta_L = \frac{p^L + u_0 - \alpha}{q^L - \alpha}$$
$$\theta_{HL} = \frac{p^H - p^L}{q^H - q^L}$$

With $\Pi_{HL} = p^L(\theta_{HL} - \theta_L) + p^H(1 - \theta_{HL})$, we have

$$\frac{\partial \Pi_{HL}}{\partial p^L} = \frac{2(p^H - p^L)}{q^H - q^L} - \frac{2p^L + u_0 - \alpha}{q^L - \alpha}$$
$$\frac{\partial \Pi_{HL}}{\partial p^H} = 1 - 2\frac{p^H - p^L}{q^H - q^L}$$

Taking first order conditions, we get

$$\theta^*_{HL} = \frac{1}{2},\ \theta^*_L = \frac{q^L + u_0 - 2\alpha}{2(q^L - \alpha)}$$
$$p^{H*} = \frac{q^H - u_0}{2},\ p^{L*} = \frac{q^L - u_0}{2}$$

As long as $0 \le \theta^*_L < \theta^*_{HL}$, that is, $\begin{cases} q^L > \alpha > u_0 \\ q^L + u_0 \ge 2\alpha \end{cases}$, the above is the firm's optimal product and pricing strategy, that is, to offer both types of products, with $p^{H*} = \frac{q^H - u_0}{2}$, $p^{L*} = \frac{q^L - u_0}{2}$.

Otherwise, the firm offers only the high-end product, with $p^{H*} = \frac{q^H - u_0}{2}$, $\theta^*_H = \frac{u_0 + q^H - 2\alpha}{2(q^H - \alpha)}$ and $\Pi^*_H = \frac{(q^H - u_0)^2}{4(q^H - \alpha)}$.[12] □

Proof. ***Proposition 1.2.***

The necessary condition for freemium to be optimal is $\frac{\partial \Pi_{HL}}{\partial p^L}|_{p^{L*}=0,p^{H*}} \le 0$. We prove this can be satisfied when

$\left[\frac{\partial V^H(\theta_{HL},D)}{\partial D} - \frac{\partial V^L(\theta_{HL},D)}{\partial D}\right]|_{p^{L*}=0,p^{H*}} \ge \left[\frac{1 + \frac{\int_{\theta_L}^{\theta_{HL}} f(\theta)d\theta}{p^H f(\theta_{HL})}\left(\frac{\partial V^H(\theta_{HL},D)}{\partial \theta_{HL}} - \frac{\partial V^L(\theta_{HL},D)}{\partial \theta_{HL}}\right)}{f(\theta_L)\frac{\partial \theta_L}{\partial p^L}}\right]|_{p^{L*}=0,p^{H*}}$,

where the right-hand side is a positive value.

Recall from the proof of Proposition 1, we have

$$\frac{\partial \Pi_{HL}}{\partial p^L} = \int_{\theta_L}^{\theta_{HL}} f(\theta)\,d\theta + p^L\left[f(\theta_{HL}) - f(\theta_L)\right]\left(\frac{\partial \theta_{HL}}{\partial p^L} - \frac{\partial \theta_L}{\partial p^L}\right) - p^H f(\theta_{HL})\frac{\partial \theta_{HL}}{\partial p^L}$$

According to the above equation, the necessary condition $\frac{\partial \Pi_{HL}}{\partial p^L}|_{p^{L*}=0,p^{H*}} \le 0$ can be satisfied when

$$\frac{\partial \theta_{HL}}{\partial p^L}|_{p^{L*}=0,p^{H*}} \ge \frac{\int_{\theta_L}^{\theta_{HL}} f(\theta)\,d\theta}{p^H f(\theta_{HL})}|_{p^{L*}=0,p^{H*}} \tag{1.6}$$

Next we determine the value of $\frac{\partial \theta_{HL}}{\partial p^L}|_{p^{L*}=0,p^{H*}}$.

[12]Note our implicit assumption is that the firm has incentive to enter the market, thus $\Pi^*_H > 0$, i.e., $q^H > \alpha$.

Recall from equation (2), we have $V^H(\theta_{HL}, D) - p^H = V^L(\theta_{HL}, D) - p^L$. By implicit function theorem, taking first derivative w.r.t. p^L on both sides, we can get

$$\frac{\partial V^H(\theta_{HL}, D)}{\partial \theta_{HL}}\frac{\partial \theta_{HL}}{\partial p^L} + \frac{\partial V^H(\theta_{HL}, D)}{\partial D}\frac{\partial D}{\partial p^L} = \frac{\partial V^L(\theta_{HL}, D)}{\partial \theta_{HL}}\frac{\partial \theta_{HL}}{\partial p^L} + \frac{\partial V^L(\theta_{HL}, D)}{\partial D}\frac{\partial D}{\partial p^L} - 1$$

Rearranging, we have

$$\frac{\partial \theta_{HL}}{\partial p^L} = \frac{\left[\frac{\partial V^H(\theta_{HL},D)}{\partial D} - \frac{\partial V^L(\theta_{HL},D)}{\partial D}\right]\left(-\frac{\partial D}{\partial p^L}\right) - 1}{\frac{\partial V^H(\theta_{HL},D)}{\partial \theta_{HL}} - \frac{\partial V^L(\theta_{HL},D)}{\partial \theta_{HL}}} \tag{1.7}$$

Substituting $\frac{\partial \theta_{HL}}{\partial p^L}$ expressed by equation (7) into the inequality (6), we have

$$\left[\frac{\left[\frac{\partial V^H(\theta_{HL},D)}{\partial D} - \frac{\partial V^L(\theta_{HL},D)}{\partial D}\right]\left(-\frac{\partial D}{\partial p^L}\right) - 1}{\frac{\partial V^H(\theta_{HL},D)}{\partial \theta_{HL}} - \frac{\partial V^L(\theta_{HL},D)}{\partial \theta_{HL}}}\right]|_{p^{L*}=0, p^{H*}} \geq \frac{\int_{\theta_L}^{\theta_{HL}} f(\theta)\, d\theta}{p^H f(\theta_{HL})}|_{p^{L*}=0, p^{H*}}$$

Rearranging, we have

$$\left[\frac{\partial V^H(\theta_{HL}, D)}{\partial D} - \frac{\partial V^L(\theta_{HL}, D)}{\partial D}\right]|_{p^{L*}=0, p^{H*}} \geq \left[\frac{1 + \frac{\int_{\theta_L}^{\theta_{HL}} f(\theta) d\theta}{p^H f(\theta_{HL})}\left(\frac{\partial V^H(\theta_{HL},D)}{\partial \theta_{HL}} - \frac{\partial V^L(\theta_{HL},D)}{\partial \theta_{HL}}\right)}{f(\theta_L)\frac{\partial \theta_L}{\partial p^L}}\right]|_{p^{L*}=0, p^{H*}}$$

Notice that the above rearrangement can be got because $\frac{\partial\left[V^H(\theta_{HL},D) - V^L(\theta_{HL},D)\right]}{\partial \theta_{HL}} > 0$ and $\frac{\partial D}{\partial p^L}|_{p^L=0} < 0$. The former is given by assumption AIV. Here we show the latter. Following the same logic in the proof of Proposition 1, when $\frac{\partial V^L(\theta,D)}{\partial D} \neq \frac{\partial V^H(\theta,D)}{\partial D}$, we still have $\frac{\partial \theta_L}{\partial p^L} > 0$. For Π_{HL}, we have $D = \int_{\theta_L}^1 f(\theta)\, d\theta$, and $\frac{\partial D}{\partial p^L}|_{p^L=0} = -f(\theta_L)\frac{\partial \theta_L}{\partial p^L} < 0$. With $\frac{\partial \theta_L}{\partial p^L} > 0$, $f(\theta) > 0$ and $\frac{\partial\left[V^H(\theta_{HL},D) - V^L(\theta_{HL},D)\right]}{\partial \theta_{HL}} > 0$ (AIV), it follows that $\left[\frac{1 + \frac{\int_{\theta_L}^{\theta_{HL}} f(\theta) d\theta}{p^H f(\theta_{HL})}\left(\frac{\partial V^H(\theta_{HL},D)}{\partial \theta_{HL}} - \frac{\partial V^L(\theta_{HL},D)}{\partial \theta_{HL}}\right)}{f(\theta_L)\frac{\partial \theta_L}{\partial p^L}}\right]|_{p^{L*}=0, p^{H*}}$ is a positive value. □

Proof. ***Corollary 1.2 and Corollary 1.3.***

According to Lemma 1, we have

$$p^L = \theta_L\left(q^L - \alpha^L\right) + \alpha^L - u_0,$$
$$p^H = \left(q^H - q^L\right)\theta_{HL} + \alpha^H + \theta_L\left(q^L - \alpha^H\right) - u_0.$$

The total profit is

$$
\begin{aligned}
\Pi_{HL}\left(\theta_{HL},\theta_{L}\right) &= p^{H}\left(1-\theta_{HL}\right)+p^{L}\left(\theta_{HL}-\theta_{L}\right) \\
&= \left[\left(q^{H}-q^{L}\right)\theta_{HL}+\alpha^{H}+\theta_{L}\left(q^{L}-\alpha^{H}\right)-u_{0}\right]\left(1-\theta_{HL}\right) \\
&+ \left[\theta_{L}\left(q^{L}-\alpha^{L}\right)+\alpha^{L}-u_{0}\right]\left(\theta_{HL}-\theta_{L}\right) \\
s.t.\ & 0\leq\theta_{L}<\theta_{HL}<1 \\
& \theta_{L}\left(q^{L}-\alpha^{L}\right)+\alpha^{L}-u_{0}\geq 0
\end{aligned}
$$

We can see that $\Pi_{HL}\left(\theta_{HL},\theta_{L}\right)$ is concave in both θ_{HL} and θ_{L}. Therefore, we can get the global optimal $\theta^{*}_{HL},\theta^{*}_{L}$ that maximize Π_{HL}, by employing the first derivatives of Π_{HL} w.r.t. θ_{HL},θ_{L}, respectively. Then p^{L*} and p^{H*} can be obtained.

$$
\frac{\partial\Pi_{HL}}{\partial\theta_{HL}} = \theta_{L}\left(\alpha^{H}-\alpha^{L}\right)+\alpha^{L}-\alpha^{H}+\left(q^{H}-q^{L}\right)\left(1-2\theta_{HL}\right)
$$

$$
\begin{aligned}
\frac{\partial\Pi_{HL}}{\partial\theta_{L}} &= \left(q^{L}-\alpha^{H}\right)\left(1-\theta_{HL}\right)+\left(q^{L}-\alpha^{L}\right)\left(\theta_{HL}-2\theta_{L}\right)-\alpha^{L}+u_{0} \\
&= \theta_{HL}\left(\alpha^{H}-\alpha^{L}\right)-2\theta_{L}\left(q^{L}-\alpha^{L}\right)+q^{L}-\alpha^{H}-\alpha^{L}+u_{0}
\end{aligned}
$$

Alternatively, we can write the marginal consumer types as $\theta_{L}=\frac{p^{L}+u_{0}-\alpha^{L}}{q^{L}-\alpha^{L}}$, $\theta_{HL}=\frac{p^{H}-p^{L}-\left(\alpha^{H}-\alpha^{L}\right)\frac{q^{L}-p^{L}-u_{0}}{q^{L}-\alpha^{L}}}{q^{H}-q^{L}}$, and the total demand is $D=1-\theta_{L}$. With $0<\theta_{L}<1$ and $0<\theta_{H}<1$, we must have $p^{L}+u_{0}<q^{L}$ and $q^{L}>\alpha^{L}$. We can express the profit as

$$
\begin{aligned}
\Pi_{HL}\left(p^{L},p^{H}\right) &= p^{L}\left(\theta_{HL}-\theta_{L}\right)+p^{H}\left(1-\theta_{HL}\right) \\
&= p^{L}\left[\frac{\left(p^{H}-p^{L}\right)\left(q^{L}-\alpha^{L}\right)-\left(\alpha^{H}-\alpha^{L}\right)\left(q^{L}-p^{L}-u_{0}\right)}{\left(q^{H}-q^{L}\right)\left(q^{L}-\alpha^{L}\right)}-\frac{p^{L}+u_{0}-\alpha^{L}}{q^{L}-\alpha^{L}}\right] \\
&+ p^{H}\left[1-\frac{\left(p^{H}-p^{L}\right)\left(q^{L}-\alpha^{L}\right)-\left(\alpha^{H}-\alpha^{L}\right)\left(q^{L}-p^{L}-u_{0}\right)}{\left(q^{H}-q^{L}\right)\left(q^{L}-\alpha^{L}\right)}\right]
\end{aligned}
$$

When $0\leq\frac{p^{L*}+u_{0}-\alpha^{L}}{q^{L}-\alpha^{L}}<\frac{p^{H*}-p^{L*}-\left(\alpha^{H}-\alpha^{L}\right)\frac{q^{L}-p^{L*}-u_{0}}{q^{L}-\alpha^{L}}}{q^{H}-q^{L}}<1$, $0<p^{L*}<p^{H*}$ and $\Pi^{*}_{HL}\geq\frac{\left(q^{H}-u_{0}\right)^{2}}{4\left(q^{H}-\alpha^{H}\right)}$, the firm would offer two products with positive prices and

$$
\theta^{*}_{L}=\frac{\left(\alpha^{H}-\alpha^{L}\right)\left(\alpha^{H}-\alpha^{L}-q^{H}+q^{L}\right)+2\left(q^{H}-q^{L}\right)\left(\alpha^{H}+\alpha^{L}-u_{0}-q^{L}\right)}{\left(\alpha^{H}-\alpha^{L}\right)^{2}-4\left(q^{H}-q^{L}\right)\left(q^{L}-\alpha^{L}\right)}
$$

$$\theta_{HL}^{*}=\frac{\left(\alpha^{H}-\alpha^{L}\right)^{2}+\alpha^{L}\left(u_{0}+2q^{H}-3q^{L}\right)+\alpha^{H}\left(q^{L}-u_{0}\right)-2q^{L}\left(q^{H}-q^{L}\right)}{\left(\alpha^{H}-\alpha^{L}\right)^{2}-4\left(q^{H}-q^{L}\right)\left(q^{L}-\alpha^{L}\right)}$$

$$\Pi_{HL}^{*}=\frac{\left(q^{H}-q^{L}\right)\left[\alpha^{L}\left(q^{H}-u_{0}\right)-\alpha^{H}\left(q^{L}-u_{0}\right)+\left(2u_{0}-q^{H}\right)q^{L}-u_{0}^{2}\right]}{\left(\alpha^{H}-\alpha^{L}\right)+4\left(\alpha^{L}-q^{L}\right)\left(q^{H}-q^{L}\right)}$$

When $\frac{\partial\Pi_{HL}}{\partial p^{L}}|_{p^{L*}=0,p^{H*}>0}\leq 0$, $0<\frac{u_0-\alpha^L}{q^L-\alpha^L}<\frac{\left(q^H-q^L\right)q^L-\alpha^L\left(q^H-2q^L+u_0\right)+\alpha^H\left(u_0-q^L\right)}{2\left(q^H-q^L\right)\left(q^L-\alpha^L\right)}<1$ and $\Pi_{HL}^{*}\geq\frac{\left(q^H-u_0\right)^2}{4\left(q^H-\alpha^H\right)}$, the firm offers two products with:

$p^{L*}=0$, $p^{H*}=\frac{q^L\left(\alpha^H-q^L\right)-\left(\alpha^H-\alpha^L\right)u_0}{2\left(q^L-\alpha^L\right)}+\frac{q^H}{2}$,

$\theta_L^*=\frac{u_0-\alpha^L}{q^L-\alpha^L}$,$\theta_{HL}^*=\frac{\left(q^H-q^L\right)q^L-\alpha^L\left(q^H-2q^L+u_0\right)+\alpha^H\left(u_0-q^L\right)}{2\left(q^H-q^L\right)\left(q^L-\alpha^L\right)}$,

$\Pi_{HL}^*=\frac{\left[\alpha^L\left(q^H-u_0\right)-q^L\left(q^H-q^L\right)+\alpha^H\left(u_0-q^L\right)\right]^2}{4\left(q^H-q^L\right)\left(q^L-\alpha^L\right)^2}$.

Otherwise the firm offers only the high-end product with $p^{H*}=\frac{q^H-u_0}{2}$,$\theta_H^*=\frac{q^H+u_0-2\alpha^H}{2\left(q^H-\alpha^H\right)}$, $\Pi_H^*=\frac{\left(q^H-u_0\right)^2}{4\left(q^H-\alpha^H\right)}$
. □

Proof. ***Concavity of profit function in q^i.***

When q^i is an endogenous decision, we prove that, if $V^i(\theta,D)$ is concave in q^i and $C\left(q^i\right)$ is convex in q^i for $i\in\{H,L\}$, the profit function is concave in q^i .

First, we look at the case where only the high-end product is provided. By Lemma 1, we have θ_H defined by $V^H(\theta_H,D)-p^H=u_0$, so $p^H=V^H(\theta_H,D)-u_0$.

$$\begin{aligned}\Pi_H&=\left[p^H-C\left(q^H\right)\right]\int_{\theta_H}^{1}f\left(\theta\right)d\theta\\&=\left[V^H(\theta_H,D)-u_0-C\left(q^i\right)\right]\int_{\theta_H}^{1}f\left(\theta\right)d\theta\end{aligned}$$

$$\frac{\partial\Pi_H}{\partial q^H}=\left[\frac{\partial V^H(\theta_H,D)}{\partial q^H}-\frac{\partial C\left(q^i\right)}{\partial q^H}\right]\int_{\theta_H}^{1}f\left(\theta\right)d\theta$$

$$\frac{\partial^2\Pi_H}{\partial\left(q^H\right)^2}=\left[\frac{\partial^2 V^H(\theta_H,D)}{\partial\left(q^H\right)^2}-\frac{\partial^2 C\left(q^i\right)}{\partial\left(q^H\right)^2}\right]\int_{\theta_H}^{1}f\left(\theta\right)d\theta$$

Therefore, when $V^i(\theta,D)$ is concave in q^i, and $C\left(q^i\right)$ is convex in q^i for $i\in\{H,L\}$, we have $\frac{\partial^2\Pi_H}{\partial\left(q^H\right)^2}<0$; thus, Π_H is concave in q^H.

Following the same logic, we can get Π_{HL} is concave in q^H and q^L. $\square$

Proof. ***Proposition 1.3.***

In conventional product line design

$$\begin{aligned}
\Pi_{HL} = (1-\theta_{HL})&\left[p^H - c\left(q^H\right)^2\right] + (\theta_{HL}-\theta_L)\left[p^L - c\left(q^L\right)^2\right] \\
s.t. p^L &= \theta_L q^L + \alpha^L(1-\theta_L) - u_0 \\
p^H &= \theta_{HL}\left(q^H - q^L\right) + \left(\alpha^H - \alpha^L\right)(1-\theta_L) + p^L \\
0 &< \theta_L < \theta_{HL} < 1
\end{aligned}$$

Because Π_{HL} is concave in q^i, with first order conditions w.r.t. q^i, we have:

$$\begin{cases} q^{H*} = \frac{\theta_{HL}}{2c} \\ q^{L*} = \frac{\theta_{HL}+\theta_L-1}{2c} \end{cases}$$

Substituting q^{H*} and q^{L*} into Π_{HL}, we can derive the optimal desicision of $\theta^*_{HL}, \theta^*_L$:

$$\begin{cases} \theta^*_{HL} = \frac{1}{15}\left(11 + 8\alpha^L c - \sqrt{1 - 64\alpha^L c + 64\left(\alpha^L\right)^2 c^2 + 60cu_0}\right) \\ \theta^*_L = \frac{1}{15}\left(7 + 16\alpha^L c - 2\sqrt{1 - 64\alpha^L c + 64\left(\alpha^L\right)^2 c^2 + 60cu_0}\right) \end{cases}$$

or

$$\begin{cases} \theta^*_{HL} = \frac{1}{15}\left(11 + 8\alpha^L c + \sqrt{1 - 64\alpha^L c + 64\left(\alpha^L\right)^2 c^2 + 60cu_0}\right) \\ \theta^*_L = \frac{1}{15}\left(7 + 16\alpha^L c + 2\sqrt{1 - 64\alpha^L c + 64\left(\alpha^L\right)^2 c^2 + 60cu_0}\right) \end{cases}$$

subject to $0 < \theta^*_L < \theta^*_{HL} < 1$. Substituting into $q^{H*}$and q^{L*}, we can get $\frac{\partial q^{L*}}{\partial \alpha^L} > 0$.

When freemium is optimal, we have

$$\begin{aligned}
\Pi_F = (1-\theta_{HL})&\left[p^H - c\left(q^H\right)^2\right] + (\theta_{HL}-\theta_L)\left[-c\left(q^L\right)^2\right] \\
s.t. p^L &= 0 = \theta_L q^L + \alpha^L(1-\theta_L) - u_0 \\
p^H &= \theta_{HL}\left(q^H - q^L\right) + \left(\alpha^H - \alpha^L\right)(1-\theta_L) \\
0 &< \theta_L < \theta_{HL} < 1
\end{aligned}$$

As Π_F is concave in q^i. With first order conditions w.r.t q^H, we have $q^{H*} = \frac{\theta_{HL}}{2c}$. Taking derivative w.r.t q^L, we have $\frac{\partial \Pi_F}{\partial q^L} < 0$ always holds.

$$\begin{cases} q^{H*} = \frac{\theta_{HL}}{2c} \\ q^{L*} = \frac{u_0 - \alpha^L(1-\theta_L)}{\theta_L} \end{cases} \tag{1.8}$$

With (8) , we have

$$\frac{\partial q^{L*}}{\partial \alpha^L} = \frac{\theta_L - 1}{\theta_L} < 0$$

□

Proof. ***Proposition 1.4.***

With $\alpha^H = \alpha^L = \alpha$, we now have

$$V^L(\theta, \alpha, D) = \theta \left(q^L + \alpha D \right),$$
$$V^H(\theta, \alpha, D) = \theta \left(q^H + \alpha D \right),$$

where D is the total demand of all offered products. As is typical in games with network effects, multiple equilibria may exist in the second stage. We seek the Nash Equilibrium that is Pareto dominant. More specifically, when network effects are intermediate, there exist multiple equilibria where all, some, or none of the consumers adopt the products. When consumers do not adopt, the products do not generate sufficient network effects and thus non-adoption becomes self-fulfilling. This coordination failure is classic in models with network effects. Clearly, the equilibrium wherein all users adopt generates (weakly) higher surplus for all parties. Thus, we select that equilibrium whenever it exists. For the proof of Proposition 4, two cases are analyzed below.

Case 1: Sell to both segments with $(q^H, \alpha, p^H), (q^L, \alpha, p^L)$

When network effects are present, the binding constraints in the firm's optimization problem continue to be the low-end consumers' IR constraint and the high-end consumers' IC constraint. The optimal prices satisfy:

$$p^L = \theta_L(q^L + \alpha) - u_0$$

$$p^H = \theta_H(q^H - q^L) + \theta_L(q^L + \alpha) - u_0$$

We bound q^L above zero in the following analysis. The optimal qualities can therefore be determined by a profit maximizing problem wherein:

$$\Pi_{HL} = \lambda\left[p^H - c\left(q^H\right)^2 - s\alpha^{\mathbf{2}}\right] + (1-\lambda)\left[p^L - c\left(q^L\right)^2 - s\alpha^{\mathbf{2}}\right] \ s.t. \ p^L{\geq}0, \ q^L{\geq}0$$

Using Lagrangian method, the optimal menu is

$q^{H*} = \frac{\theta_H}{2c}$

$$q^{L*} = \begin{cases} \frac{\theta_L - \theta_H\lambda}{2c(1-\lambda)} & , \lambda\theta_H\theta_L{\leq}\theta_L^2 \ \& \ \frac{s\left[\theta_H\theta_L\lambda + 2cu_0(1-\lambda)\right]}{\theta_L^2(c-c\lambda+s)} - 1 \leq 0 \\ 0 & , \lambda\theta_H\theta_L > 2su_0 > \theta_L^2 \\ 0 & , \lambda\theta_H\theta_L > \theta_L^2 \geq 2su_0 \\ \frac{2su_0\theta_L - \theta_H\lambda}{2\theta_L(c-c\lambda+s)} & , 2su_0 > \lambda\theta_H\theta_L \& 2su_0 > \theta_L^2 \& \frac{s\left[\theta_H\theta_L\lambda + 2cu_0(1-\lambda)\right]}{c-c\lambda+\theta_L^2 s} - 1 > 0 \end{cases}$$

$$\alpha^* = \begin{cases} \frac{\theta_L}{2s} & , \lambda\theta_H\theta_L{\leq}\theta_L^2 \ \& \ \frac{s\left[\theta_H\theta_L\lambda + 2cu_0(1-\lambda)\right]}{\theta_L^2(c-c\lambda+s)} - 1 \leq 0 \\ u_0 & , \lambda\theta_H\theta_L > 2su_0 > \theta_L^2 \\ \frac{\theta_L}{2s} & , \lambda\theta_H\theta_L > \theta_L^2 \geq 2su_0 \\ \frac{\theta_H\theta_L\lambda + 2cu_0(1-\lambda)}{2\theta_L(c-c\lambda+s)} & , 2su_0 > \lambda\theta_H\theta_L \& 2su_0 > \theta_L^2 \& \frac{s\left[\theta_H\theta_L\lambda + 2cu_0(1-\lambda)\right]}{c-c\lambda+\theta_L^2 s} - 1 > 0 \end{cases}$$

$$p^{L*} = \begin{cases} \frac{\theta_L\left(\theta_L - \theta_H\lambda\right)}{2c(1-\lambda)} + \frac{\theta_L^2}{2s} - u_0 & , \lambda\theta_H\theta_L{\leq}\theta_L^2 \ \& \ \frac{s\left[\theta_H\theta_L\lambda + 2cu_0(1-\lambda)\right]}{\theta_L^2(c-c\lambda+s)} - 1 \leq 0 \\ 0 & , \lambda\theta_H\theta_L > 2su_0 > \theta_L^2 \\ \frac{\theta_L^2}{2s} - u_0 & , \lambda\theta_H\theta_L > \theta_L^2 \geq 2su_0 \\ 0 & , 2su_0 > \lambda\theta_H\theta_L \& 2su_0 > \theta_L^2 \& \frac{s\left[\theta_H\theta_L\lambda + 2cu_0(1-\lambda)\right]}{c-c\lambda+\theta_L^2 s} - 1 > 0 \end{cases}$$

$$p^{H*} = \begin{cases} \frac{\theta_H^2 + \theta_L^2 - \theta_H\theta_L(1+\lambda)}{2c(1-\lambda)} + \frac{1}{2s} - u_0 & , \lambda\theta_H\theta_L{\leq}\theta_L^2 \ \& \ \frac{s\left[\theta_H\theta_L\lambda + 2cu_0(1-\lambda)\right]}{\theta_L^2(c-c\lambda+s)} - 1 \leq 0 \\ \frac{\theta_H^2}{2c} & , \lambda\theta_H\theta_L > 2su_0 > \theta_L^2 \\ \frac{\theta_H^2}{2c} + \frac{\theta_L^2}{2s} - u_0 & , \lambda\theta_H\theta_L > \theta_L^2 \geq 2su_0 \\ \theta_H\left(\frac{\theta_H}{2c} - \frac{2su_0 - \theta_H\theta_L\lambda}{2\theta_L(c-c\lambda+s)}\right) & , 2su_0 > \lambda\theta_H\theta_L \& 2su_0 > \theta_L^2 \& \frac{s\left[\theta_H\theta_L\lambda + 2cu_0(1-\lambda)\right]}{c-c\lambda+\theta_L^2 s} - 1 > 0 \end{cases}$$

$$\Pi_{HL}^* = \begin{cases} \frac{\theta_L^2 + \theta_H^2\lambda - 2\theta_H\theta_L\lambda}{4c(1-\lambda)} + \frac{\theta_L^2}{4s} - u_0 & , \lambda\theta_H\theta_L{\leq}\theta_L^2 \ \& \ \frac{s\left[\theta_H\theta_L\lambda + 2cu_0(1-\lambda)\right]}{\theta_L^2(c-c\lambda+s)} - 1 \leq 0 \\ \frac{\theta_H^2\lambda}{4c} - \frac{su_0^2}{\theta_L^2} & , \lambda\theta_H\theta_L > 2su_0 > \theta_L^2 \\ \frac{\theta_H^2\lambda}{4c} + \frac{\theta_L^2}{4s} - u_0 & , \lambda\theta_H\theta_L > \theta_L^2 \geq 2su_0 \\ \frac{\theta_H^2\theta_L^2\lambda(c+s) - 4\theta_H\theta_L scu_0\lambda - 4sc^2u_0^2(1-\lambda)}{4\theta_L^2 c(c+s-c\lambda)} & , 2su_0 > \lambda\theta_H\theta_L \& 2su_0 > \theta_L^2 \& \frac{s\left[\theta_H\theta_L\lambda + 2cu_0(1-\lambda)\right]}{c-c\lambda+\theta_L^2 s} - 1 > 0 \end{cases}$$

Case 2: Sell to θ_H-segment only with(q, α, p)

Binding condition $p = \theta_H(q + \alpha\lambda) - u_0$. Firm's profit is given by $\Pi_H = \lambda(p - cq^2 - s\alpha^2)$. Solving the optimization problem, we obtain:

$$p^* = \frac{\theta_H^2}{2c} + \frac{\lambda^2\theta_H^2}{2s} - u_0, \ q^* = \frac{\theta_H}{2c}, \ \alpha^* = \frac{\lambda\theta_H}{2s}$$

$$\Pi_H^* = \lambda\left(\frac{\theta_H^2}{4c} + \frac{\lambda^2\theta_H^2}{4s} - u_0\right)$$

Freemium is optimal if and only if (IFF) the conditions $\Pi_{HL}^* \geq \Pi_H^*,\ p^{L*} = 0$ are satisfied.

Below, we prove the above conditions cannot hold simultaneously. According to the results obtained by using the Lagrangian method, we discuss by parameter ranges where p^{L*} may possibly be zero.

(1)When $\lambda\theta_H\theta_L > 2su_0 > \theta_L^2$.

We have

$$\begin{aligned}
\Pi_H^* - \Pi_{HL}^* &= \lambda\left(\frac{\theta_H^2\lambda^2}{4s} - u_0\right) + \frac{su_0^2}{\theta_L^2} \\
&> \lambda\left(u_0\frac{su_0}{\theta_L^2} - u_0\right) + \frac{u_0}{2} \\
&> u_0\left(\frac{1}{2} - \frac{\lambda}{2}\right) > 0
\end{aligned}$$

Therefore, $\Pi_{HL}^* < \Pi_H^*$ always holds when $\lambda\theta_H\theta_L > 2su_0 > \theta_L^2$; thus, the firm prefers to offer only the high-end product, and freemium cannot emerge.

(2) When $2su_0 > \lambda\theta_H\theta_L$, $2su_0 > \theta_L^2$ and $\frac{s[\theta_H\theta_L\lambda + 2cu_0(1-\lambda)]}{c - c\lambda + \theta_L^2 s} - 1 > 0$.

Following the same logic as in (1), we prove $\Pi_H^* > \Pi_{HL}^*$ always holds as long as the high-type consumer finds the low-end product worth trying at price zero, i.e., $\theta_L(q^L + \lambda\alpha) - u_0 \geq 0$.

$$\begin{aligned}
\Pi_{HL}^* &= \frac{\theta_H^2\theta_L^2\lambda(c+s) - 4\theta_H\theta_L scu_0\lambda - 4sc^2u_0^2(1-\lambda)}{4\theta_L^2 c(c+s-c\lambda)} \\
&= \frac{\lambda\theta_H^2}{4c} + \frac{\theta_H^2\lambda^2}{4(c+s-c\lambda)} - \frac{su_0[\theta_H\theta_L\lambda + cu_0(1-\lambda)]}{\theta_L^2(c+s-c\lambda)}
\end{aligned}$$

$$\Pi_H^* = \frac{\lambda\theta_H^2}{4c} + \lambda\left(\frac{\lambda^2\theta_H^2}{4s} - u_0\right)$$

Let $F_{HL} = \frac{\theta_H^2\lambda^2}{4(c+s-c\lambda)} - \frac{su_0[\theta_H\theta_L\lambda + cu_0(1-\lambda)]}{\theta_L^2(c+s-c\lambda)}$, $F_H = \lambda\left(\frac{\lambda^2\theta_H^2}{4s} - u_0\right)$. The IFF condition for $\Pi_H^* > \Pi_{HL}^*$ is $F_H > F_{HL}$. We proceed by proving that, under the conditon $p^{L*} < 0$, $F_H > F_{HL(max)}$ holds; thus, $F_H > F_{HL}$ and $\Pi_H^* > \Pi_{HL}^*$.

Below we prove F_{HL} is decreasing in c; thus, $\underset{c>0}{sup}F_{HL} = \underset{c\to 0}{lim}F_{HL} = F_{HL}|_{c=0} = \frac{\theta_H^2\lambda^2}{4s} - \frac{u_0\theta_H\lambda}{\theta_L}$. Therefore, a sufficient condition for $F_H > F_{HL}$ is $F_H > \underset{c>0}{sup}F_{HL}$, or equivalently, $F_H > F_{HL}|_{c=0}$. (Notice F_H is not a function of c.)

$$\frac{\partial F_{HL}}{\partial c} = -\frac{\theta_H^2\theta_L^2\lambda(1-\lambda)}{4\theta_L^2(c+s-c\lambda)} - \frac{4su_0(su_0-\theta_H\theta_L\lambda)(1-\lambda)}{4\theta_L^2(c+s-c\lambda)}$$
$$= -\frac{(1-\lambda)[4su_0(su_0-\theta_H\theta_L\lambda)+\theta_H^2\theta_L^2\lambda]}{4\theta_L^2(c-c\lambda+s)}$$

We have $4su_0(su_0-\theta_H\theta_L\lambda)+\theta_H^2\theta_L^2\lambda > 4\frac{\lambda\theta_H\theta_L}{2}(\frac{\lambda\theta_H\theta_L}{2}-\theta_H\theta_L\lambda)+\theta_H^2\theta_L^2\lambda = -\theta_H^2\theta_L^2\lambda^2+\theta_H^2\theta_L^2\lambda > 0$, hence $\frac{\partial F_{HL}}{\partial c} < 0$, and F_{HL} is decreasing in c.

Next we prove $F_H > F_{HL}|_{c=0}$. When c approaches 0, the condition $\frac{s[\theta_H\theta_L\lambda+2cu_0(1-\lambda)]}{\theta_L^2(c-c\lambda+s)} - 1 > 0$ implies $\lambda > \frac{\theta_L}{\theta_H}$. We first prove $F_{HL}|_{c=0} < 0$ for all $\lambda > \frac{\theta_L}{\theta_H}$.

At $\lambda > \frac{\theta_L}{\theta_H}$, we have

$$F_{HL}\big|_{c=0,\lambda=\frac{\theta_L}{\theta_H}} = \frac{\theta_L^2}{4s} - u_0 < \frac{2su_0}{4s} - u_0 = -\frac{u_0}{2} \leq 0$$

and

$$\frac{\partial F_{HL}|_{c=0}}{\partial\lambda} = \frac{\theta_H^2\lambda}{2s} - \frac{u_0\theta_H}{\theta_L} < \frac{2su_0/\theta_L}{2s}\theta_H - \frac{u_0\theta_H}{\theta_L} = 0$$

Therefore, when $\lambda > \frac{\theta_L}{\theta_H}$, we always have $F_{HL}|_{c=0} < 0$.

We also have

$$\begin{aligned} F_H &= \lambda\left(\frac{\lambda^2\theta_H^2}{4s} - u_0\right) \\ &> \lambda\left(\frac{\lambda^2\theta_H^2}{4s} - u_0\frac{\theta_H}{\theta_L}\lambda\right) \\ &= \lambda F_{HL}|_{c=0} \\ &> F_{HL}|_{c=0} \text{ (since } F_{HL}|_{c=0} < 0) \end{aligned}$$

So we have proved $F_H > F_{HL}|_{c=0}$; hence, $F_H > F_{HL}$ is also proved. □

Proof. ***Proposition 1.5.***

The proof follows similar logic as in Proposition 4. With asymmetric network effects, consumer valuations of the products are:

$$V^L(\theta,\alpha^L,D) = \theta(q^L+\alpha^L D),$$
$$V^H(\theta,\alpha^H,D) = \theta(q^H+\alpha^H D).$$

Case 1: Sell to both segments with $(q^H, \alpha^H, p^H), (q^L, \alpha^L, p^L)$.

The optimal prices satisfy:

$$p^L = \theta_L(q^L + \alpha^L) - u_0$$
$$p^H = \theta_H(q^H - q^L + \alpha^H - \alpha^L) + \theta_L(q^L + \alpha^L) - u_0$$

The optimal qualities can therefore be determined by a profit maximizing problem wherein:

$$\Pi_{HL} == \lambda\left[p^H - c\left(q^H\right)^2 - s\left(\alpha^H\right)^2\right] + (1-\lambda)\left[p^L - c\left(q^L\right)^2 - s\left(\alpha^L\right)^2\right] \ s.t.\ p^L{\geq}0,\ q^L{\geq}0$$

The optimal quality levels remain the same as in the no-network-effects scenario, namely

$q^{H*} = \frac{\theta_H}{2c}$

$$q^{L*} = \begin{cases} \frac{\theta_L - \theta_H\lambda}{2c(1-\lambda)} & , \frac{2csu_0(1-\lambda)}{\theta_L^2(c+s)} + \frac{\lambda\theta_H}{\theta_L} - 1 \leq 0 \\ \frac{su_0}{\theta_L(c+s)} & , \frac{2csu_0(1-\lambda)}{\theta_L^2(c+s)} + \frac{\lambda\theta_H}{\theta_L} - 1 > 0 \end{cases}$$

$\alpha^{H*} = \frac{\theta_H}{2s}$

$$\alpha^{L*} = \begin{cases} \frac{\theta_L - \theta_H\lambda}{2s(1-\lambda)} & , \frac{2csu_0(1-\lambda)}{\theta_L^2(c+s)} + \frac{\lambda\theta_H}{\theta_L} - 1 \leq 0 \\ \frac{cu_0}{\theta_L(c+s)} & , \frac{2csu_0(1-\lambda)}{\theta_L^2(c+s)} + \frac{\lambda\theta_H}{\theta_L} - 1 > 0 \end{cases}$$

$$p^{L*} = \begin{cases} \frac{\theta_L(\theta_L - \theta_H\lambda)}{2(1-\lambda)}\left(\frac{1}{c} + \frac{1}{s}\right) - u_0 & , \frac{2csu_0(1-\lambda)}{\theta_L^2(c+s)} + \frac{\lambda\theta_H}{\theta_L} - 1 \leq 0 \\ 0 & , \frac{2csu_0(1-\lambda)}{\theta_L^2(c+s)} + \frac{\lambda\theta_H}{\theta_L} - 1 > 0 \end{cases}$$

$$p^{H*} = \begin{cases} \frac{\theta_H^2 + \theta_L^2 - (1+\lambda)\theta_L\theta_H}{2(1-\lambda)}\left(\frac{1}{c} + \frac{1}{s}\right) - u_0 & , \frac{2csu_0(1-\lambda)}{\theta_L^2(c+s)} + \frac{\lambda\theta_H}{\theta_L} - 1 \leq 0 \\ \frac{\theta_H^2}{2}\left(\frac{1}{c} + \frac{1}{s}\right) - \frac{u_0\theta_H}{\theta_L} & , \frac{2csu_0(1-\lambda)}{\theta_L^2(c+s)} + \frac{\lambda\theta_H}{\theta_L} - 1 > 0 \end{cases}$$

The optimal profit is:

$$\Pi_{HL}^* = \begin{cases} \frac{\theta_H^2 + \theta_L^2\lambda - 2\lambda\theta_L\theta_H}{4(1-\lambda)}\left(\frac{1}{c} + \frac{1}{s}\right) - u_0 & , \frac{2csu_0(1-\lambda)}{\theta_L^2(c+s)} + \frac{\lambda\theta_H}{\theta_L} - 1 \leq 0 \\ \frac{\theta_H^2\lambda(c+s)}{4cs} - \frac{\theta_H\lambda u_0}{\theta_L} - \frac{csu_0^2(1-\lambda)}{\theta_L^2(c+s)} & , \frac{2csu_0(1-\lambda)}{\theta_L^2(c+s)} + \frac{\lambda\theta_H}{\theta_L} - 1 > 0 \end{cases}$$

Case 2: Sell to θ_H-segment only

Here the situation is the same as in case 2 in the symmetric network effects senario, where $p = \theta_H(q + \alpha\lambda) - u_0$ and $\Pi_H = \lambda(p - cq^2 - s\alpha^2)$:

$$q^* = \frac{\theta_H}{2c},\ \alpha^* = \frac{\lambda\theta_H}{2s},$$
$$\Pi_H^* = \lambda\left(\frac{\theta_H^2}{4c} + \frac{\lambda^2\theta_H^2}{4s} - u_0\right)$$

For freemium to be optimal, the conditions $\Pi_{HL}^* \geq \Pi_H^*,\ p^{L*} = 0$ have to be satisfied.

Under asymmetric network effects, freemium equilibrium exists, when (s, c, u_0, λ) satisfy the conditions below:

$$\frac{\theta_H^2}{4s}(\lambda - \lambda^3) + \lambda u_0 \left(1 - \frac{\theta_H}{\theta_L}\right) - \frac{cs(1-\lambda)u_0^2}{\theta_L^2(c+s)} \geq 0$$
$$\frac{{\theta_H}^{\mathbf{2}}\lambda(c+s)}{4sc} - \frac{cs(1-\lambda)u_0^2}{\theta_L^2(c+s)} - \frac{\theta_H \lambda u_0}{\theta_L} > 0$$
$$\frac{2csu_0\,(1-\lambda)}{\theta_L^2(c+s)} + \frac{\lambda\theta_H}{\theta_L} - 1 > 0$$

It can be seen that the above inequalities define a non-empty set. Corollary 4 follows straightforwardly according to the above results. □

Appendix B

In the analysis in section 1.5.2, we alluded to the possibility that the assumption "the high-type consumers have positive valuation for the low-end product at price zero" may not hold. This means that the highest-type consumers will not find the low-end product worth trying even if it is offered at zero price, which is very unlikely for successful freemium products. In this section, we provide a detailed analysis for this exceptional case and show that even if the assumption is violated, our results still hold true unless the cost function has some special form. With some special cost function and the assumption violated (i.e., $V^L(\theta_H, \alpha, \lambda) < 0$), raising the low-end product's price above zero will violate the high type's IR instead of IC constraint, then it is possible that freemium would emerge under uniform network effects.

As in the baseline model, the firm can either serve the high type consumers with menu (p, q, α_1), getting profit Π; or serve both high and low segments, with menu $(p^H, q^H, \alpha_2), (p^L, q^L, \alpha_2)$, getting profit Π_{HL}. Freemium is a special case of the second strategy, where $p^L = 0$ and profit Π_F. We focus on necessary conditions for freemium to be optimal.

Let $C_q(q)$, $C_\alpha(\alpha)$ be the marginal cost for offering one product of quality q and network benefit α.

(1)When $\lambda > \frac{\theta_L}{\theta_H}$.

By checking only p and p^H, we can see that the firm prefers offering only the high-end product to adopting freemium strategy. More specifically, consider p^{H*}:

$$
\begin{aligned}
p^{H*} &= \theta_H(q^{H*} - q^{L*}) + \theta_L(q_L^* + \alpha_2^*) - u_0 \\
&< \theta_H(q^{H*} - q^{L*}) + \theta_H q_L^* + \theta_L \alpha_2^* - u_0 \\
&= \theta_H q^{H*} + \theta_L \alpha_2^* - u_0
\end{aligned}
$$

Now let's consider p, where we let the firm sets $\alpha_1 = \alpha_2^*$:

$$
\begin{aligned}
p &= \theta_H(q^* + \lambda\alpha_2^*) - u_0 \ (IR\ constraint) \\
&= \theta_H q^{H*} + \theta_H \lambda \alpha_2^* - u_0 \ (since\ q^* = q^{H*}) \\
&> \theta_H q^{H*} + \theta_L \alpha_2^* - u_0 \ (since\ \lambda > \frac{\theta_L}{\theta_H}) \\
&> p^{H*}
\end{aligned}
$$

Therefore, compared to serving only the high-type consumers, it is never optimal to pursue the freemium strategy, because the firm will get less profit on the high end while subsidizing the low-end segment. In fact, when selling only to the high-type customers, the firm can even increase profit by adopting an optimal α_1^*, making $\Pi^* > \Pi_{HL}^*$ always hold under $\lambda > \frac{\theta_L}{\theta_H}$.

(2)When $\lambda \leq \frac{\theta_L}{\theta_H}$.

First, consider serving only the high end. We have:

$$
\begin{aligned}
p &= \theta_H(q + \lambda\alpha_1) - u_0 \\
\Pi &= \lambda[p - C_q(q) - C_\alpha(\alpha_1)] \\
&= \lambda[\theta_H(q + \lambda\alpha_1) - u_0 - C_q(q) - C_\alpha(\alpha_1)]
\end{aligned}
$$

To maximize Π, we have:

$$
\begin{aligned}
C_q{'}(q^*) &= \theta_H \\
C_\alpha{'}(\alpha_1^*) &= \lambda\theta_H
\end{aligned}
$$

for $\Pi^* > 0$, we have$u_0 < \theta_H(q^* + \lambda\alpha_1^*) - C_q(q^*) - C_\alpha(\alpha_1^*)$.

Then, when the firm serves both segments, we have:

$$p^H = \theta_H(q^H - q^L) + \theta_L(q^L + \alpha_2) - u_0$$
$$p^L = \theta_L(q^L + \alpha_2) - u_0$$
$$\Pi_{HL} = \lambda[p^H - C_q(q^H)] + (1-\lambda)[p^L - C_q(q^L)] - C_\alpha(\alpha_2)$$

To maximize Π_{HL}, we have:

$$C_q{}'(q^{H*}) = \theta_H$$
$$C_q{}'(q^{L*}) = \frac{\theta_L - \lambda\theta_H}{1-\lambda}$$
$$C_\alpha{}'(\alpha_2^*) = \theta_L$$

Notice the necessary condition for freemium is $p^{L*} \leq 0$, where $p^{L*} = \theta_L(q^{L*} + \alpha_2^*) - u_0$. Under this condition, we check whether the firm would like to induce more cost (by offering higher-than-optimal $q^L + \alpha_2$ to make $p^L = 0$) to have the low-type consumers on board. Supposing the firm adopts freemium, we need to compare Π^* and Π_F, where the firm set q^L and α_2 such that $p_L = \theta_L(q^L + \alpha_2) - u_0 = 0$ and $q^{H*} > q^L \geq q^{L*}$. Notice that $q^* = q^{H*}$, and q^{H*} is always equal to the efficient quality.

Assuming a convex cost function, with $u_0 < \theta_H(q^* + \lambda\alpha_1^*) - C_q(q^*) - C_\alpha(\alpha_1^*)$, we have:

$$\begin{aligned}
\Delta\Pi &= \Pi_F^* - \Pi^* \\
&= \lambda(u_0 - \alpha_1^*\lambda\theta_H - \theta_H q^L) - (1-\lambda)C_q(q^L) + \lambda C_\alpha(\alpha_1^*) - C_\alpha(\alpha_2) \\
&< \lambda(\theta_H q^{H*} - C_q(q^*) - \theta_H q^L) - (1-\lambda)C_q(q^L) - C_\alpha(\alpha_2) \\
&< \lambda[C_q{}'(q^{H*})(q^{H*} - q^L) - C_q(q^{H*})] - (1-\lambda)C_q(q^L)
\end{aligned}$$

We can see that $\Delta\Pi < 0$ unless $C_q(\cdot)$ is very steep and skewed towards zero, specifically, $C_q\prime(q^{H*}) > \left(\frac{1-\lambda}{\lambda}\right)\frac{C_q(q^L)}{q^{H*}-q^L} + \frac{C_q(q^{H*})}{q^{H*}-q^L}$.

Above we analyzed the case where network effect is endogenous. As can be easily seen from the analysis, exactly the same conclusion can be reached for the case where network effects are exogenously given, no matter how large α is.

Chapter 2

Are Hype News Corrected or Amplified? The Oz Effect in Healthcare

2.1. Introduction

As Reuters reported,[1] health information is one of the most frequently sought information on the Internet. On average, 53% of American people search for health information online.[2] While publicly available health information can sometimes help us better manage our health, alleviate concerns about our wellness, or even avoid some hospital visits, it can also be the source for incorrect or misleading information. In 2016, the sensational tragedy of Zexi Wei in China drew unprecedented public attention to the issue of credibility of online healthcare information.[3] Baidu.com was scolded as an accessory to murder because it did not pull misleading medical information that recommended unproven methods to treat a rare form of cancer from its search results. Given the enormous growth of medical information online, especially from less credible sources, there may be a serious risk of erroneous information or exaggerated information influencing medical treatments. What exacerbates the problem is that deceptive information has been propelled to popularity through exposure in mainstream media supported by trusted

---

[1] Reuters. Consumer-targeted internet investment: online strategies to improve patient care and product positioning. Reuters

[2] Pew Internet American Life Project. Internet visits soaring. Health Management Technology 2003; 24: 2–8.

[3] Wikipedia (2017). *Death of Wei Zexi*. Retrieved from https://en.wikipedia.org/wiki/Death_of_Wei_Zexi

spokespeople. For example, Dr. Oz, known for his talk show, *The Dr. Oz Show*,[4] referred to an ingredient called Green Coffee Bean Extract (GCBE) as a "magic weight-loss cure" and a "miracle pill" that can burn fat fast without diet or lifestyle change. Because of the show's popularity, sales of the products that contain GCBE skyrocketed shortly afterward. However, a study published in the British Medical Journal (Korownyk et al. 2014) on the effectiveness of Oz's medical advice found that only 46 percent of his recommendations had any scientific backing or rationale. The study showed that 39 percent had no supporting scientific evidence, while the remaining 15 percentage points went directly against scientific evidence. Concerned about the negative impact of the Dr. Oz Show, in 2014, Senate's consumer protection panel grilled Dr. Oz about his promotion of supposed weight-loss cures like GCBE. During the Senate hearing, Dr. tried to justify the usage of these forceful words by claiming that, "…people act on emotion and how they feel, so a main principle in building our scripts is to ellicit a visceral, emotional reaction from the viewer." This type of messages are not meant to let consumers know both helpful and harmful facts about the product, but are hyped and designed to steer consumers towards extremes. Such biased information is especially harmful in the case of serious public health challenges such as obesity that seem to defy most treatments.

Although reputable media organizations have often played a role of "gatekeeper" to protect consumers from misleading information in hype news, this role has been challenged in the internet age. Product reviews, despite their authenticity, may come from consumers who cannot identify fraud. Research articles, on the other hand, may not generate helpful findings quickly

[4] We use "*The Dr. Oz Show*" and "the Oz Show" interchangeably throughout the paper.

enough. Therefore, it remains unclear whether public information sources could correct misleading health information or not. This study fills this research gap by studying how the biased information from celebrity doctors affect the supply of public information for over-the-counter (OTC) weight loss products. For OTC healthcare products, consumers' purchasing decisions are especially prone to be affected by the publicly available information, because consumers are free to choose what information to believe and do not need a physician's prescription. Weight loss product is a good representative of OTC healthcare products. In the US, more than 70% of adults aged 20 and over are overweight, including obesity,[5] and more than 30% of people who made weight loss attempts used OTC weight loss products (Eisenberg, 2008). And globally, the revenue for weight loss and weight management market is expected to increase from $15.9 billion in 2016 to $22.9 billion by 2025. Given the importance of the obesity issue and the enormous market value (Khan et al., 2016), we choose to focus on weight loss products in this study.

As we analyze the cascade of public health information, we collect data from several sources, in order to cover the entire spectrum. Public information sources are different in credibility and audience coverage. But all of them play an indispensable role in affecting consumer's healthcare choices (e.g., Hu & Sundar, 2010; Bates et al., 2006). According to potentially different coverage and credibility levels, we analyze public health information from the following four sources:

1. Peer-reviewed scientific research articles. Compared to other sources, research articles are highly credible, because of the serious review process and expertise of the medical researchers.

---

[5] Center for Disease Control and Prevention (2013-2014). Retrieved from https://www.cdc.gov/nchs/fastats/obesity-overweight.htm

However, they are harder to reach mass consumers due to the narrow outlet and the abstruse content.

2. User-generated product reviews. On the one hand, customer reviews can convey first-hand evaluation of the healthcare product/service; thus, they can be very credible and reliable. On the other hand, reviews can become problematic because a customer may not know how to give an unbiased evaluation. Moreover, reviews can be deceptive if they are manipulated by the product/service provider.

3. Genuine news articles written by professional journalists or reporters. The articles included in our analysis are from legitimate and genuine news agencies, as opposed to fake news. News articles may not be highly credible because news reporters or journalists do not always write with evidence that can stand the test of time, as they may lack the expertise and time to mine the truth from all information they collect.

4. TV shows launched or participated in by doctors. A good representative of this source is *The Dr. Oz Show*. Despite the success of his show, he has been criticized by physicians and government officials for giving non-scientific advice.[6] There are also other TV shows (e.g., The Doctors) that recommend weight-loss products but focus mostly on other medical issues. Given *The Dr. Oz Show* supplies the majority of this type of information according to the data, we categorize all these shows and call them the Oz Show for ease of reference. We also refer to the effect of the show as the "Oz Effect" through this paper.

---

[6] In 2014, The Federal Trade Commission filed a lawsuit against several products he peddled on his show, and he was scolded during the congressional hearing.

For each information source, we examine various dimensions of language features contained in each piece of information by employing the state-of-the-art deep-learning-based natural language processing (NLP) techniques. Specifically, we address three research questions: First, how does the hype news from *The Dr. Oz Show* affect the information generation from other sources (i.e., news articles, UGC, and research articles)? Second, does the information from the other sources amplify or mitigate the "bias" in the hype news? Lastly, how does the market respond to the Oz Show from the perspectives of consumer search and product supply? To answer these three questions, we employ two types of research designs: Regression Discountinuity in Time (Lee and Lemieuxa 2010; Hausman and Rapson 2017) and Synthetic Control (Abadie et al. 2010, 2015).

Our findings consist of several components. First, we find that the hype news from the Oz Show leads to increasing consumer search for the recommended ingredients,[7] more publicly available information generated from news articles, as well as higher price of the products containing the concerned ingredient. This means that consumers try to look for more information after listening to hype news, news articles try to meet this increasing demand for information, and firms respond by increasing the price. Second, a noteworthy finding is that, news articles from genuine outlets are acting as a magnifier rather than rectifier. We find, by analyzing the language features used in the news content, that the vast majority of news articles are written with higher sentiment, no significant change in emotions, and little correction for the ingredients peddled by Dr. Oz on the show. Regarding the more credible information source, namely scholarly journals,

[7] We will use "ingredient" and "(product) category" interchangeably in this paper.

only one out of thousands of peer-reviewed articles directly correct what Dr. Oz said on the talk show. Some consumers correct the biased information in UGC but they are overwhelmed by much more content supporting the hype news or remaining neutral. This implies that government intervention in the health information domain is crucial in order to protect consumers from misleading information.

The Dr. Oz Show serves as an example to demonstrate a general implication that, healthcare information from celebrity doctors may initiate a media-hype (Vasterman 2005; Zuckerman 2003), creating news waves on the recommended healthcare products. While exaggerated information causes attention, consumers rely on searching through credible sources to make decisions. However, what makes the problem more severe is that, when hype news happens, it seems to drive real news in the same direction. While Google, Baidu, and Facebook are fighting fake news in healthcare domain, our finding suggests that even legitimate and genuine information sources can hurt consumers. News articles respond to the hype news by propagaging and amplifying the hype news, leaving consumers vulnerable to misleading health information. Research articles either react too slowly or do not respond to the biased information. Though user-generated content (UGC) provides some correction, it may be dominated by the larger amount of content that supports the misleading information. In fact, the recent action by Instagram of hiding anti-vaccine misinformation[8], which is identified as false information by World Health Organization, the Centers for Disease Control, and similar organizations, further

[8] https://www.theverge.com/2019/5/9/18553821/instagram-anti-vax-vaccines-hashtag-blocking-misinformation-hoaxes

demonstrates that, only after government or reputable public sectors identify the misleading information, other information intermediaries can get a clear guide to battle it.

In summary, this paper makes a substantive contribution to provide concrete empirical evidence on how a biased and hyped health information source affects subsequent information generation from other sources. From a public policy perspective, the results have important implications on how to potentially regulate media content in the healthcare domain in order to protect consumers from misleading health information. The regulators may want to generate more bias correction articles for public access after hype news happens, because other more credible sources could not seem to offer effective corrections for the biased information.

We organize the rest of the paper as follows. In Section 2.2, the related literature and the contribution of the present paper are discussed. Section 2.3 describes the data. Section 2.4 introduces how we extract useful features from the text data using deep learning in natural language processing. Section 2.5 discusses the results. Finally, in Section 2.6, we conclude with managerial and public policy implications, as well as limitations.

2.2. Literature review

This paper relates to the literature of information diffusion in marketing (e.g., Tirunillai and Tellis 2017) and sociology (e.g., Goel et al. 2012). Tirunillai and Tellis (2017) study how offline TV advertising affects online user-generated content, including consumer reviews and blogs. Goel et al. (2012) examine the tree structure of online diffusion network. In comparison, we study the healthcare information cascade along multiple channels (i.e., news, UGC, and research), given all are typically accessible to consumers thus influence consumer decisions. We are especially interested in how the hyped information is passed along different channels, and whether it can be corrected by more credible sources (i.e., genuine news articles and research papers). Our results about the Oz Effect on information intensity is also consistent with Goel et al. (2012)'s finding that the bulk of information adoption often takes place within one degree of a few dominant individuals.

Our paper is related to the literature on media bias (e.g., Gentzkow and Shapiro 2006, 2010). Gentzkow and Shapiro (2006) take a theoretical approach to analyze the supply-side incentive of generating biased news, such as reputation concerns. Gentzkow and Shapiro (2010) examine what drives media slant empirically, taking into account of firm's profit maximization goal and consumer's subjective ideology. In comparison, our study focuses on the healthcare market, instead of the domain of politics. Moreover, we focus on how the bias cascades by separately examining the language features and intensity of information flow, instead of how the bias was originated.

This paper also relates to the striving literature on the impact of fake news in various fields (e.g., Rao and Wang 2017; Allcott and Gentzkow 2017; Vosoughi et al. 2018; Friggeri et al. 2014). This paper differentiates from the fake news literature by studying how news information from legitimate sources can work against consumers.

This paper also relates to the communication literature on media hype (e.g., Vasterman 2005), which describes the phenomenon of self-inflating media coverage on one specific story or topic. Vasterman (2005) uses a case study to examine the dynamics of media-hype. As a more specific study on hype in medical news, Zuckerman (2003) conducts three case studies of how companies shape news coverage of medical products. Goel et al. (2012) study the online information diffusion structures with seven examples. They find that most trends on social media form only after mentions from a few dominant individuals or news outlet amplification. In contrast, our paper presents empirical evidence of media hype on healthcare products—*The Dr. Oz Show* serving as the key event and news articles forming consonant news waves.

Our paper is also related to the pharmaceutical marketing literature. For example, Ching et al. (2016) investigates the impact of publicity on demand for Anti-Cholesterol drugs. Chintagunta et al. (2009) focus on the learning process of doctors. Kalra et al. (2011) study the impact of negative and positive media coverage on physicians' beliefs about the quality of a prescription-based diabetes drug. In sum, all three papers above study the effect of news coverage on prescription choices, which are made by doctors rather than patients. In comparison, our study focuses on the supply side of public health information, and we speak mainly to over-the-counter healthcare products, where consumers heavily rely on public healthcare-related information to make decisions.

More broadly speaking, this paper is related to the literature on applying natural language processing to marketing (e.g., Lee and Bradlow 2011; Tirunillai and Tellis 2012; Lee et al. 2017; Liu et al. 2017). Instead of solely relying on human coding or feature engineering, we also employ the state-of-the-art deep learning approach (LeCun et al. 2015) to content code all textual information from different sources.

2.3. Data

To collect publicly available information generated from various channels, we combine several datasets: user generated data (i.e., customer reviews), news articles, research articles, and scripts of popular TV shows. In this section, we provide an overview of our data sources and the data collection approach.

2.3.1. Products Information

The first dataset we use is the Amazon product data from 1996 to 2014,[9] including product reviews and product-level metadata. We focus on the category of health and personal care and further narrow down to the weight loss subcategory. We replace the cross-sectional data on price with the time-series price information for each product, enabled by the API developed by Keepa.com. For each product, we collect the monthly advertisement spending data from Ad$pender database at the brand level.[10] In total, there are 6119 weight loss products and 150,731 consumer-generated reviews. Table 2.1 shows the summary statistics at the product level. Notice that the ranking is measured in the entire health and personal care category. Table 2.1 shows some examples of weight loss products in our sample.

From the products' titles, we manually extract the key ingredients. For example, the ingredient of the first product listed in Table 2.2 is garcinia cambogia, whereas that of the second product is raspberry ketones. For each ingredient, we check whether it has been recommended on *The Dr.*

---

[9] Made available by Julian McAuley, UCSD. http://cseweb.ucsd.edu/~jmcauley/. Our goal is to study how hype news affect information cascade, so we restrict the time window of our study to be before the FTC hearing (June 17th, 2014) to keep the study context clean. Therefore, the data we collected and analyzed for all information sources (news, UGC, research articles) are all before June 1st, 2014.

[10] The user manual is available here: http://products.kantarmediana.com/documents/AdSpenderManual.pdf

Oz Show. Among all the products, 1864 contain the ingredients mentioned by Dr. Oz. We list all ingredients that are identified from the product's titles.

Table 2.1. Product-level Descriptive Statistics

Variable	Mean	Std.	Min	Max
Number of Reviews	7.664	26.439	1	847
Price	23.704	24.149	0.010	1025.380
Average Rating	3.490	1.353	1	5
Average Ranking	104765.500	116826	55	1205512
Ads. Spending (000)	6.243	111.648	0	4771.600

Table 2.2. Examples of Product Titles

1	Garcinia Cambogia Extract by NewLife Botanicals
2	NatureWise Raspberry Ketones Plus+ Weight Loss Supplement and Appetite Suppressant
3	Lipozene Diet Pills - Maximum Strength Fat Loss Formula - 1500mg - 30 Capsules
4	Power Pops-hoodia Weightloss Lollipops-30ct Variety Pack
5	nuYou Labs Green Coffee Bean Extract with GCA Chlorogenic Acid - Highly Effective Natural Weight Loss Diet Supplement
6	One XS Weight Loss Pills (X-Strength) Prescription Grade Diet Pill. No Prescription Needed. Fast Proven Results. Weight Loss Guarantee
7	Molecular Research Labs Diet Supplement, Garcinia Cambogia Extract, 750 mg, 60 Count
8	NOW Foods Liver Detoxifier and Regenerator, 90 Capsules
9	Eden Pond Ketones Liquid Diet Drops Best Fat Burner Weight Loss That Works, Raspberry, 2 Fluid Ounce
10	Trimspa x32 Rapid Release Weight Loss 70 Capsules

- **Ingredients mentioned on the Dr. Oz Show**: Garcinia cambogia, green coffee bean, raspberry ketone, saffron extract, forskolin, safflower oil, moringa, glucomannan, chitosan, 7-keto.[11]

[11] We also detected ingredients not mentioned on the Dr. Oz Show but are also natural ingredients: e.g., Caralluma, citrimax, sesamin.

In addition to the products that contain the ingredients listed above, some products contain no clear ingredients in their titles (e.g., the last title in Table 2.2), while some do not work in a similar way as the "Oz products."[12] Therefore, when collecting information generated from various channels at the ingredient level, for these products with no mention of ingredient in the product title, their consumer reviews will not be included.

Table 2.3. Number of Products, Reviews, News, and Research Articles at Ingredient Level

Ingredient	(1) #product	(2) #review	(3) #news	(4) #research
Garcinia Cambogia	740	42933	522	116
Green Coffee Bean	425	19557	491	73
Raspberry Ketones	431	14365	430	14
saffron extract	63	1390	42	2141
forskolin	80	1188	40	1157
safflower oil	8	210	479	67
moringa	15	161	54	306
glucomannan	19	469	132	251
chitosan	20	167	184	17
7-keto	12	32	22	1191

2.3.2. User-generated Content

The user-generated data used for our analysis is consumer reviews from Amazon.com. As shown in Table 2.1, on average, there are 7.664 reviews posted for each product listing. In Table 2.3, we list the total number of products and consumer reviews at the ingredient level in the first and second columns.

[12] The active ingredients in the "Oz products" are all natural products. However, other products in the weight-loss category are made of artificial chemical compounds or use physical mechanisms to assist weight loss. These products include Diuretic, Diurex, Enema, Hydroxycut, Nuphedrine, and Ornithine.

2.3.3. Scripts of the Oz Show

For each of the identified ingredients from Amazon products' title, we searched the corresponding Dr. Oz Show that recommended the focal ingredient. For each episode, we extracted the script for further textual analysis. Each ingredient was featured in one episode, except that Green Coffee Bean were covered in two episodes.

2.3.4. News articles

We collect news articles from LexisNexis, one of the major databases for news articles, covering 15,000 news resources in the database. In our study, we define news article as all those included in LexisNexis in the following categories: Newswires & Press Releases, Newspapers, and News Transcripts.

Similar to the approach used in Ching et al. (2016), when collecting news articles, we search for articles that contain the ingredient name (e.g., garcinia cambogia). To make sure the article is about our interested topic (i.e., weight loss), we keep only articles that contain the keyword "weight loss" or "lose weight" or "fat". Every ingredient received some news coverage. In Table 2.3, we list the total number of news articles collected at the ingredient level in the third column.

2.3.5. Peer-reviewed research articles

We collect peer-reviewed research articles from the ProQuest Central database. There are well over 10,000 scholarly journals indexed in ProQuest Central, covering all major subject areas, including business, health and medical, social sciences, arts and humanities, education, science and technology, and religion. The search and collection strategy is similar to that for news articles, except that we only keep the articles in the category "Scholarly Journal" of ProQuest

Central. Please see the appendix B for the full list of journals from which we collect the related research articles. The fourth column of Table 2.3 reports the number of peer-reviewed research articles collected at the ingredient level.

2.4. Text Processing

To process the unstructured textual data introduced in 2.3, we go through the following procedure:

1. Use Convolutional Neural Network (CNN), a deep learning based NLP model, to identify three content features.

2. Extract sentiment using traditional NLP.

These two steps aim to extract all textual features, detailed in Section 2.4.1. The models used to extract these features are introduced in Section 2.4.2 and 2.4.3. Specifically, Section 2.4.2 describes the CNN model (step 1) while Section 2.4.3 presents the traditional NLP models (step 2). After processing the text data, we will incorporate the content features as variable in the empirical analysis in Section 2.5.

2.4.1. Textual Features to Extract

In total, we examine four textual features. Please see Table 2.4 for the definition of all textual features.

For the first three textual features listed in Table 2.4, we use CNN to extract them from each information source. Before describing CNN in 2.4.2, we first explain these features. As the purpose of the hype news from *The Dr. Oz Show* is to induce visceral and emotional reaction from the viewer, we measure whether each article (review) contains any emotion, either positive or negative. Morever, we are also curious to see whether information from other sources can act as a rectifier so that the biased information can be corrected. To extract these features, CNN is used for the following two reasons. First of all, CNN has outstanding performance on NLP tasks (e.g., Liu et al., 2017; Timoshenko & Houser, 2017), and it fits our need as a scalable supervised prediction technique to detect whether a text document contains a specific feature. Second, there are no well-established traditional NLP tools to content code these features.

Table 2.4. Description of Textual Features

Variables	Methods	Description
Bias Correction	CNN	dummy, equals 1 if the text corrects the message of Oz Show.
Positive Emotion	CNN	dummy, equals 1 if positive emotion appears in the text.
Negative Emotion	CNN	dummy, equals 1 if negative emotion appears in the text.
Sentiment	Trad.NLP	the measure of sentiment of the focal textual information.
Intensity	Summation	the measure of how many units of the textual information.

Note: "Trad.NLP" means tranditional NLP methods.

In addition to these three features, sentiment is also extracted. We choose this feature because it is also a good measure of how positive the subsequent information related to the concerned ingredient. We use traditional NLP tools to extract sentiment because they are supported by large-size external corpus and have shown good performance and robustness. We explain each of the feature extraction tasks in detail in Section 2.4.3.

The last feature we consider is intensity. Now we describe how we define and construct the intensity measure. As explained in Table 2.4, intensity is a measure of the density of the textual information in each period. We measure intensity with a simple approach by counting the frequency of all existing pieces of information at each time point. For example, the intensity for "research" of ingredient j in period t would be 2 if two new research articles about j were published in period t. For each information source, we use "cumulative sum" for the intensity measure and "average" for other features,[13] because "sum" captures the intensity of the information, while "average" provides an overall measure for all other language features.

2.4.2. Informational Content Extraction with CNN

As mentioned before, we use CNN to extract the first three features (i.e., bias correction, positive emotion, negative emotion) listed in Table 2.4. Specifically, we follow two steps to label each piece of textual information, which could be a product review, a newspaper article, a research article, or a script of an episode of *The Dr. Oz Show*. First, we hire workers through Amazon Mechanical Turk and tag 3,000 messages for a variety of textual contents. Subsequently, using the labeled contents, we train a CNN model to content code the full set of messages (more than 160,000 messages). Our CNN consists of four layers, as shown in Figure 0.1 (e.g., Kim, 2014; Liu et al., 2017; Timoshenko & Hauser, 2017). We briefly describe each layer of the CNN as follows.

- Layer 1: Word embedding.

The first layer is the word embedding or word vectors. Following the popular method to improve performance without a large supervised training set, we initialize word vectors with those obtained from an unsupervised neural language model. That is, the publicly available word2vec

---

[13] Both the "sum" and "average" for each language feature in period t are calculated across all the text documents published within period t. We also conducted robustness checks by considering not only period t, but also the past n periods, where $n \in \{1,2,3,4,5\}$. In other words, the "sum" and "average" are calculated across all the text documents published within period $t, t-1, \ldots, t-n$. The results remain largely consistent.

vectors trained on 100 billion words from Google News. Each vector has a dimension of 300 and was trained using the continuous bag-of-words architecture (Mikolov et al., 2013). New words are randomly initialized. As displayed in Figure 0.1, the first layer is the representation of the sentence, with each word represented by a 300-dimensional vector. With n being the total number of words in the text, the representation matrix is of dimension $n\times 300$. The i-th word is denoted as $\mathbf{v_i}$.

- Layer 2: Convolutional layer.

The convolutional layer applies convolutional operations with varying filters to the sentence representation in the first layer. The filter can be denoted as a vector $\mathbf{w} \in R^{1\times dh}$, which corresponds to a concatenation of all rows in a matrix from the second layer as shown in **Error! Reference source not found.**. h is the size of the filter, and d is the dimension of the word embedding (i.e., 300). In our model, three different filter sizes are implemented (e.g., 3, 4, 5). The feature map for each filter with size h is a vector of the outputs of the convolutional operation, that is,

$$\mathbf{c} = [c_1, c_2, \dots, c_{n-h+1}]$$
$$c_i = \sigma(\mathbf{w} \cdot \mathbf{v_{i:i+h-1}} + b)$$

where $\sigma(\cdot)$ is a non-linear activation function, and we use $\sigma(x) = max\{x, 0\}$. $\mathbf{w}$ is the vector of linear weights and b is the bias (i.e., intercept), both of which are to be estimated. $\mathbf{v_{i:i+h-1}}$ is a concatenation of the vectors representing words i to $i+h-1$; therefore, it is of dimension $dh\times 1$.

Layer 3: Pooling layer.

The pooling layer aims to transform the feature maps to a lower-dimensional vector, so to get the most salient textual information. The output is specified as

$$\mathbf{p} = [p_1, p_2, \dots, p_{mk}]$$

$$p_j = max\{c_1, c_2, \dots, c_{n-h+1}\}$$

where k is the total number of filters, and p_j corresponds to the output resulting from the filter of size h. We use 128 filters for each filter size h (h =3, 4, 5 in our network architecture). Therefore, there are in total $k = 128 \times 3 = 384$ filters.

Layer 4: Softmax layer

The last layer in CNN is the softmax layer. This final layer takes the output of the pooling layer (i.e., $\mathbf{p}$) as input, and outputs the probabilistic prediction of whether a feature is contained in this text. Therefore, the output is a binary result y, which equals 1 if the text is classified as containing the feature under examination. The softmax specification is

$$y = softmax(\mathbf{W} \cdot \mathbf{p} + b)$$

where the weights $\mathbf{W}$ and bias b are to be calibrated through training the CNN.

Figure 0.1. Architecture of Convolutional Neural Network for Sentence Classification

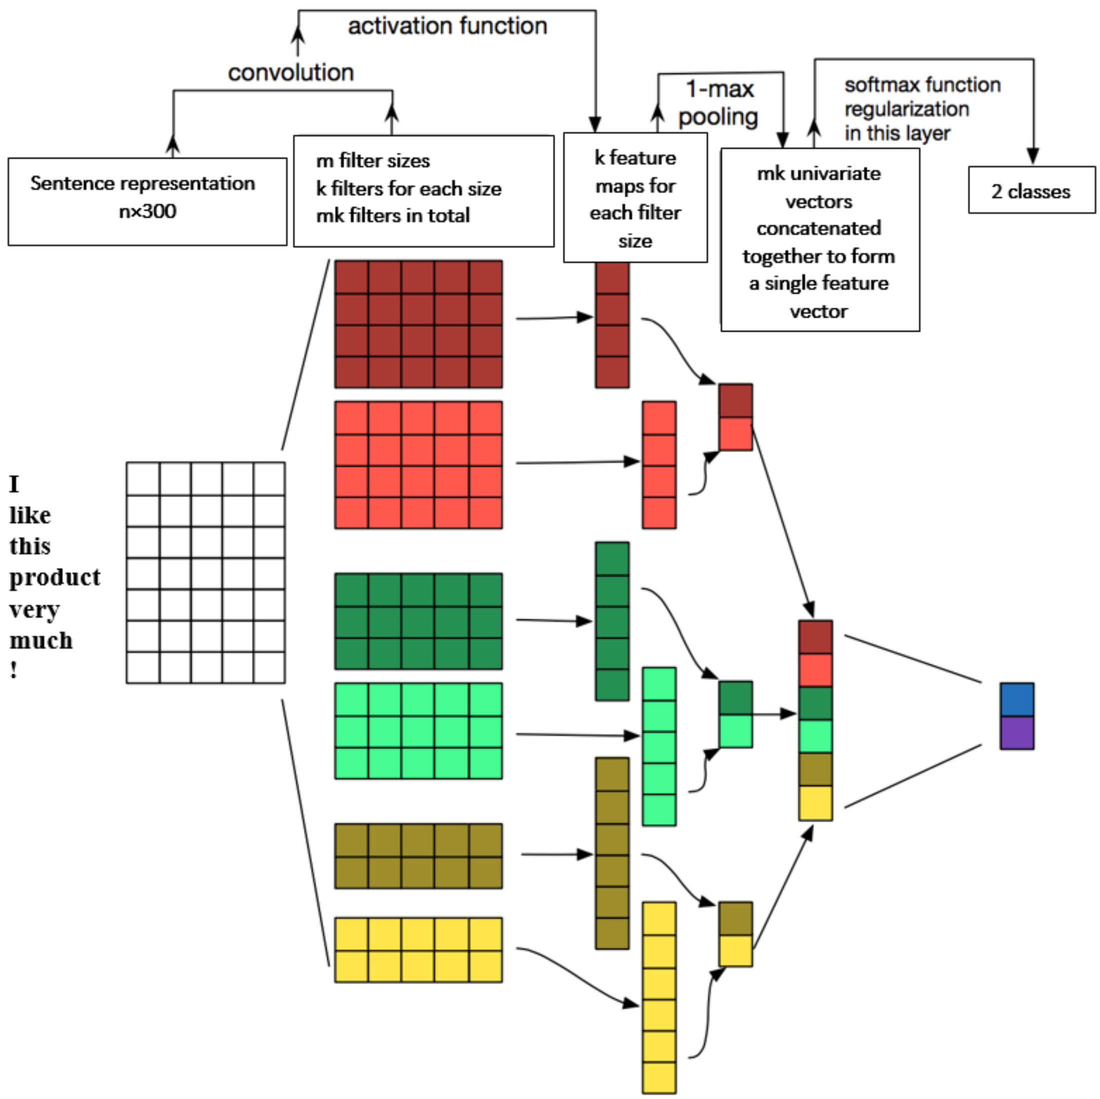

Note: in this figure, k=2, m=3, and the three filter sizes are 2, 3, 4; in our application, k=128, m=3, the filter sizes are 3, 4, 5. The figure is adapted from Figure 1 of "A Sensitivity Analysis of (and Practitioners' Guide to) Convolutional Neural Networks for Sentence Classification," by Zhang and Wallace (2016).

We use a mini-batch of size 64 while training the CNN. Following the rule-of-thumb, we set the drop-out rate as 0.5 in order to help prevent overfitting. We employ 10-fold cross validation to train the CNN. Three common criteria (i.e., precision, recall, accuracy) are used to evaluate the performance (Lee & Bradlow, 2011).[14] The CNN classifier's performance on a test sample of 1000 observations is shown in Table 2.5. Table 2.6 shows some examples of the classified text for each content feature.

Table 2.5. Performance on Content Coding

	Bias Correction	Positive Emotion	Negative Emotion
Precision (%)	87.8	72.3	80.8
Recall (%)	65.5	58.2	66.7
Accuracy (%)	88.7	79.0	85.9

Table 2.6. Classified Text for the Three Content Features Using CNN

Feature	Example Text Classified as Containing the Feature
Bias Correction	Contrary to Dr. Oz's claim, the miracle cure isn't really a miracle at all. Obesity experts are concerned over the validity of the study.
Positive Emotion	Great Product!!! If you have not tried this yet, then try it today. So glad I found it! thanks!
Negative Emotion	I have tried several kinds of this pill and none of them has done anything for me. I am very disappointed.

In addition to the three content features extracted using CNN, we use traditional NLP techniques to extract sentiment. Specifically, we use the library TextBlob for sentiment analysis. TextBlob is a high-level library built on top of the NLTK library[15] (D'Andrea et al., 2015). As we pass text to create a TextBlob object, the TextBlob library performs the following processing over text: tokenize the text, i.e., split words from body of text; remove stopwords from the tokens; POS (part of speech) tagging of the tokens and select only significant features/tokens like adjectives,

---

[14] $Precision = \frac{true\ positive}{true\ positive + false\ positive}$, $Recall = \frac{true\ positive}{true\ positive + false\ negative}$, $Accuracy = \frac{true\ prediction}{test\ sample\ size}$.

[15] For the latest version, please see https://pypi.python.org/pypi/textblob

adverbs, etc.; pass the tokens to a sentiment classifier which classifies the text sentiment as positive, negative, or neutral by assigning it a polarity score between -1 to 1.

2.4.3. Descriptive Analysis of the Extracted Text Features

Next, we provide some descriptive results of the text processing applied to different information sources. The summary statistics of the extracted language features are reported in Table 2.7.

Table 2.7. Summary Statistics of Language Features

News (Count: 2396)				
	Bias. Corr.	Neg. Emotion	Pos. Emotion	Sentiment
Mean	0.076	0.037	0.276	0.113
Std.	0.266	0.188	0.447	0.092
min	0	0	0	-0.433
50%	0	0	0	0.109
max	1	1	1	0.600
Research (Count: 5333)				
	Bias. Corr.	Neg. Emotion	Pos. Emotion	Sentiment
Mean	0.166	0.006	0.072	0.066
Std.	0.372	0.076	0.259	0.093
min	0	0	0	-0.250
50%	0	0	0	0.064
max	1	1	1	0.550
Reviews (Count: 150744)				
	Bias. Corr.	Neg. Emotion	Pos. Emotion	Sentiment
Mean	0.241	0.080	0.252	0.166
Std.	0.428	0.272	0.434	0.214
min	0	0	0	-1
50%	0	0	0	0.16
max	1	1	1	1

(a) Bias correction: Review>News≈Research.

Consumer reviews provides most correction for the hype news. Both news and research contain very little correcting information, with research ranking lowest (only one article).

(b) Positive emotion: News≈Review>Research.

Journalists or news reporters tend to use some positive emotion in the news content, similarly for consumer reviews. In comparison, research articles are more evidence-based and written in an emotionless way.

(c) Negative emotion: Review>News>Research.

We find that reviews contain the highest portion of content with negative emotion, followed by news and research articles. Consumers are more likely to express emotions after experiencing a bad product or services.

(d) Sentiment: Review>News>Research.

On an average, customer reviews contain the most positive messages, followed by news articles. The mean sentiment in research articles is the lowest and close to zero, meaning that research is almost neutral. More detailes are shown in Section 2.5.

2.5. Analysis and Results

In this section, we examine the three key research questions. First, how does the hype news from *The Dr. Oz Show* affect the information generation from other sources (i.e., news articles, UGC, and research articles)? Second, does the information from the other sources amplify or mitigate the "bias" in the hype news? Lastly, how does the market respond to the Oz Show from the perspectives of consumer search and product supply?

To answer these questions, we use two types of research designs: Regression Discountinuity in Time (RDIT) (Lee and Lemieuxa 2010; Hausman and Rapson 2017) and Synthetic Control (Abadie et al. 2010, 2015). We first describe the methods and notations.

RDIT

The estimation equation for RDIT is as follows.

$$Y_{it} = show_{it} \cdot \tau + \sum_{n=1}^{3} \delta_n t^n + \vec{\beta}\overrightarrow{X_{it}} + \xi_i + season_t + \varepsilon_{it}$$

where Y_{it} is the outcome variable of interest, which could be one of the following: information intensity, language features (postive and negative emotion, sentiment), and market response (price, product supply). For each outcome variable Y_{it}, we estimate the Oz effect on it separately. The results for each outcome variable are discussed sequentially in the rest of this section. $show_{it}$ is the treatment variable, where $show_{it} = 1$ if period t is *after* the Oz Show broadcast time for ingredient i, otherwise $show_{it} = 0$. Our key parameter of interest is τ, measuring the effect of the Oz Show.

$\overrightarrow{X_{it}}$ is the vector of characteristics for ingredicent i in period t. For news, research articles, search interest, supply of product, and price trend, $\overrightarrow{X_{it}}$ includes the advertisment spending. For UGC, it includes advertisement spending, average price, and average sales ranking within category i. We account for time-varying factors by including time polynomials $\sum_{n=1}^{3} \delta_n t^n$ up to order 3, subject to Akaike's criterion. ξ_i is the ingredient-level fixed effect which absorbs other unobserved characteristics. We also include season fixed effect, $season_t$, to account for the season-specific shocks. ε_{it} denotes the idiosyncratic shock.

Synthetic Control

Following Abadie et al. (2010, 2015), let Y_{it} represent the outcome variable (e.g., intensity, sentiment) for ingredient i at time t. Let Y_{it}^N represent the outcomes of the ingredients in the absence of the intervention, and let Y_{it}^I represent the outcomes of the focal ingredient that received the intervention. The net effect of intervention at any given time period is the gap, α_{it} , which is the difference between the treated brand and the counterfactual or synthetic brand, namely

$$\alpha_{it} = Y_{it}^I - Y_{it}^N$$

where $1 \leq t \leq T$. According to Abadie et al. (2010), the treated unit's outcome can be calculated using a convex combination of the untreated units (i.e., the synthetic control). Let $\overrightarrow{\mathrm{W}} = (w_1, \ldots, w_I)$ be the weight vector of the units. Let the treated unit be $i = 0$. With these weights, the synthetic control estimator is

$$\hat{\alpha}_{0t} = Y_{0t} - \sum_{i=1}^{I} w_i^* Y_{it}$$

w_i^* are the optimal weights to be chozen in order to minimize the difference between the characteristics of the treated unit and the synthetic control during the pre-intervention period.

We now describe the sampling strategy for the Sythetic Control method. To construct the pool of synthetic control, we first restrict the focus on ingredients that were featured on *The Dr. Oz*

Show. This is to account for unobserved heterogeneity that relates to the selection process of the Oz show. There could be some specific reasons (e.g., firm-side product strategies) that we do not observe but explain why certain ingredients were featured in the show. By restricting to only ingredients ever recommended on the Oz Show, the endogeneity issue caused by such unobserved common reasons can be eliminated. The air date for each ingredient is listed in Table 2.8 chronologically. Essentially, the control pool for an ingredient i includes all other ingredients featured later than i. In order to test the gap between the treated and the control with statistical power, we need long enough post-treatment periods, as well as a control pool that consists of reasonaly large number of units so that we can create a synthetic control to match the treated unit's pre-intervention patterns as close as possible. With this in mind, we examine the following five ingredients for the Synthetic Control analysis: Chitosan, Safflower Oil, Raspberry Ketones, 7-Keto, and Forskolin. For each of these five ingredients, we can construct a control pool containing at least four units and obtain at least eight post-treatment periods (weeks). The pairs of treatement and control pools are shown in Table 2.9. Take Chitosan as an example. The control pool for Chitosan include Green Coffee Bean, Moringa, Saffron Extract, and Garcinia Cambogia. The treatment date is 1/3/12, and the time onward up to 4/1/12 is the post-treatment period.

We define one period as one week. We choose 26 periods (half a year) of data as the preintervention period to construct the synthetic unit, which is the closest representation of the focal ingredients. During the post-intervention period, we keep the outcomes up to the end of the data sample period (i.e., June 2014) for presenting the results.

Table 2.8. Date of Mention for Each Ingredient

	Oz Show Air Date
Chitosan	1/3/12
Safflower oil	1/3/12
Forskolin	2/1/12
Raspberry Ketones	2/6/12
7-Keto	2/10/12
glucomannan	2/21/12
Green Coffee Bean	4/1/12
Moringa	4/13/12
Saffron Extract	4/17/12
Garcinia Cambogia	10/29/12

Table 2.9. Treatment vs. Control Pairs

Treated	Control Pool	
Safflower Oil/Chitosan	Green Coffee Bean/ Moringa / Saffron Extract (3 months post-treatment)	Garcinia Cambogia (10 months post-treatment)
7-Keto/Forskolin/Raspberry Ketones	Green Coffee Bean/ Moringa / Saffron Extract (2 months post-treatment)	Garcinia Cambogia (9 months post-treatment)

We use two types of variables as our predictor variables for constructing the synthetic control: (1) Past trend of the focal outcome variables (e.g., number of articles/reviews, text features), (2) Category characteristics (i.e., advertisement spending, average price, online rating).

2.5.1. How do different information channels react to hype news from the Dr. Oz show?

In this section, we explore how the hype news from the show affected the supply of public information from other sources. In 2.5.1.1, we answer the first research question: "How does the information intensity change due to the hype news?" In 2.5.1.2, we answer the second research question: "Do other sources of information amplify or mitigate the messages?" In 2.5.1.3, we provide robustness checks using the synthetic control method. And in section 2.5.1.4, we provide possible explanations for the findings.

2.5.1.1. *How does the information intensity change?*

- News articles

We show results obtained from both RDIT and Synthetic Control analysis for the news articles. For ease of exhibition, we report the RDIT results only for the other informaton sources. Both sets of results are qualitatively consistent.

Table 2.10 reports the RDIT results for Oz Effect on information intensity of the three sources. Column 1 of Table 2.10 shows the results for the Oz Effect on news article generated, with the bandwidth of two quarters (half a year) on both sides of the treatment time. Robustness tests with various bandwidths are reported in Appendix D. Figure 0.2 shows the visual demonstration of the results for the Green Coffee Bean ingredient which is representative of many other ingredients, with four different bandwidths. We will use Green Coffee Bean as a running example throughout the rest of the paper for demenstration purposes. The results show that the Oz Effect on the number of news articles generated is significantly postive. News articles from legitimate news agencies respond to the Oz Show by providing more coverage of the ingredient concerned afterwards.

Figure 0.2. Information Intensity Change over Time (e.g., Green Coffee Bean)

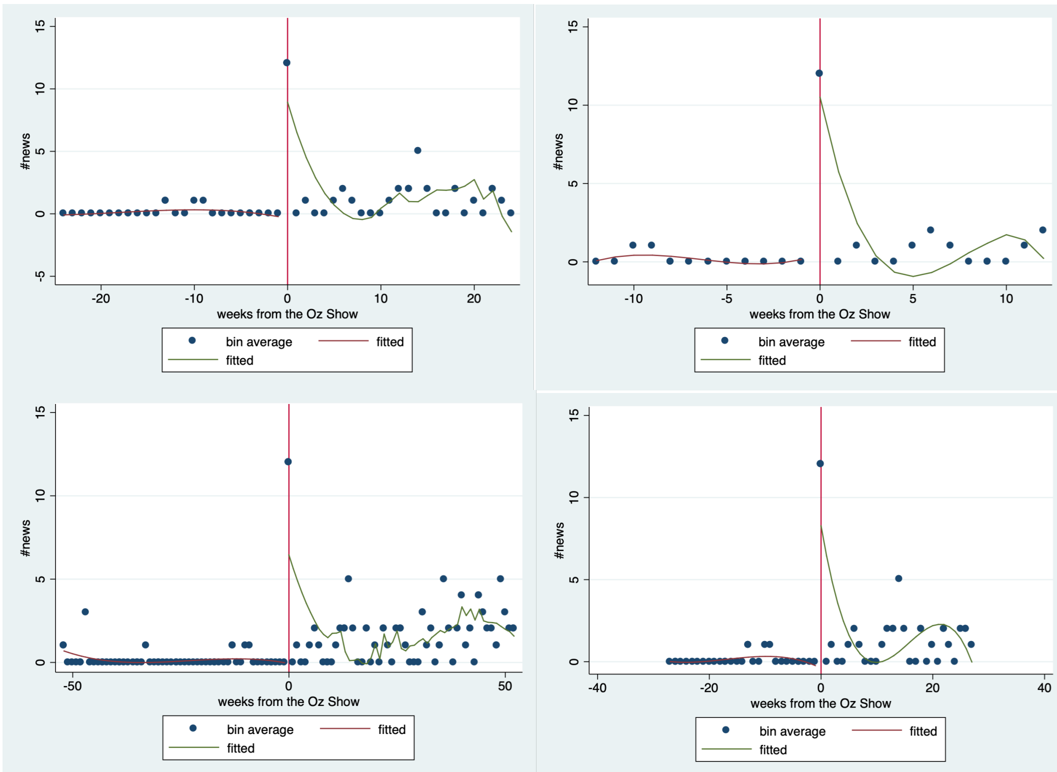

- Consumer reviews

According to the RDIT results shown in the second column of Table 2.10, the Dr. Oz Show does not have any significant causal effect on the generation of consumer reviews. It implies that consumers may not base their decisions soly on what Dr. Oz recommended.

- Research articles

The analysis of research article generation is done at monthly level.[16] In the third column of **Error! Reference source not found.**, we can see there is no significant effect of the Oz Show on the intensity of research articles. This is unsurprising due to the long publishing cycle required for scholarly journals.

Table 2.10. Oz Effect on Information Intensity

VARIABLES	(1) #News	(2) #Reviews	(3) #Research
After Oz	0.953**	7.398	0.139
	(2.53)	(0.37)	(0.06)
Ads.	0	0.134***	0.039***
	(0.03)	(6.57)	(5.96)
Price		1.366***	
		(3.02)	
Ranking		-9.418**	
		(-2.23)	
Ingredient FE	Yes	Yes	Yes
Season FE	Yes	Yes	Yes
Observations	595	665	488
Adjusted R2	0.111	0.584	0.64

Note: t-statistics in parentheses; The results are for a bandwidth of two quarters before and after the Oz Show; all with 3rd order time polynomials; *** p<0.01, ** p<0.05, * p<0.1

2.5.1.2. Does other sources of information amplify or mitigate the messages?

We have analyzed the Oz Effect on information intensity for news, UGC, and research. Next, we examine how the information content from each channel changes if a product is mentioned on the Oz Show. We are most interested in understanding whether the other information channels

[16] For research articles with only year information, we assign month from one to twelve randomly.

provide consumers a balanced view regarding the concerned weight-loss products. To achieve that, we study the Oz Effect on four aspects of the content separately: bias correction, sentiment, positive emotion, and negative emotion. **Error! Reference source not found.** shows the findings about bias correction by different channels.[17] **Error! Reference source not found.** shows the Oz Effect on sentiment and emotion.[18] Below, we organize the results by each channel: 1) news articles, 2) consumer reviews, and 3) research articles.

Table 2.11. Correction Rate by News Articles and UGC

	(1) News		(2) UGC	
	%cite Oz	%correction	%cite Oz	%correction
Garcinia Cambogia	0.4920	0.0159	0.0504	0.01522
Green Coffee Bean	0.5046	0.0162	0.0834	0.0284
Raspberry Ketones	0.5363	0.0094	0.0648	0.0211
Chitosan	0.0179	0.0	0.0229	0.0076
Forskolin	0.4762	0.0074	0.1032	0.0288
7-keto	0.0588	0.0	0.1	0.0667
Glucomannan	0.0676	0.0	0.0401	0.0094
Moringa	0.0	0.0	0.0062	0.0
Saffron Extract	0.275	0.0032	0.0927	0.0419
Safflower Oil	0.0076	0.0	0.0732	0.0098

1) News Articles

Now we present the impact of Oz Show on news articles.

<u>*Bias Correction*</u>

[17] As explained later, research articles contains near-zero correction within the study time window, therefore it is not listed in the table.
[18] The effect on negative emotion is valid only for UGC thus is shown separately in Table 2.13.

The first column of Table 2.11 reports the percentage of articles that cited Dr. Oz or his talk show, among all the news articles generated after the Oz Show. We can see quite a few categories have a high citation rate: Garcinia Cambogia, Green Coffee Beans, Raspberry Ketones, Forskolin, and Saffron Extract. However, among all news articles across these categories, less than 2% of them provide correction for the hyped language used in the Oz Show. This tells us that consumers are receiving biased information carried over by genuine news articles.

Table 2.12. Oz Effect on Sentiment and Positive Emotion

	Outcome variable: sentiment			Outcome variable: positive emotion		
	(1) News	(2) Reviews	(3) Research	(4) News	(5) Reviews	(6) Research
After Oz	0.009**	0.131***	0.025	0.005	-0.215***	-0.041
	(2.20)	(2.72)	(0.48)	(-0.32)	(-2.83)	(-0.55)
Ads.	0.000**	-0.000	0.000	0.000***	0.000	0.000
	-2.53	(-1.20)	(1.59)	(5.47)	(0.84)	(1.18)
Price		0.003**			0.004**	
		(2.37)			(2.47)	
Ranking		0.015			0.037**	
		(1.45)			(2.32)	
Ingredient FE	Yes	Yes	Yes	Yes	Yes	Yes
Season FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	595	665	130	595	665	130
Adjusted R^2	0.123	0.104	0.391	0.137	0.236	0.094

Note: t-statistics in parentheses; Bandwidth is two quarters before and after the Oz Show; all with 3rd order time polynomials; *** p<0.01, ** p<0.05, * p<0.1

<u>*Emotion*</u>

As shown in column (4) of Table 2.12, the effect of the Oz Show on positive emotion in news articles is not significant. For negative emotion, very few news articles contain negative emotions, mainly due to the requirement for professional journalistic writing. So, there is a close-to-zero variation in this language feature for news articles over a long period of time. As a result, we cannot derive any insight from it.

<u>*Sentiment*</u>

Figure 0.3. Oz Effect on Sentiment of News Articles (e.g., Green Coffee Bean)

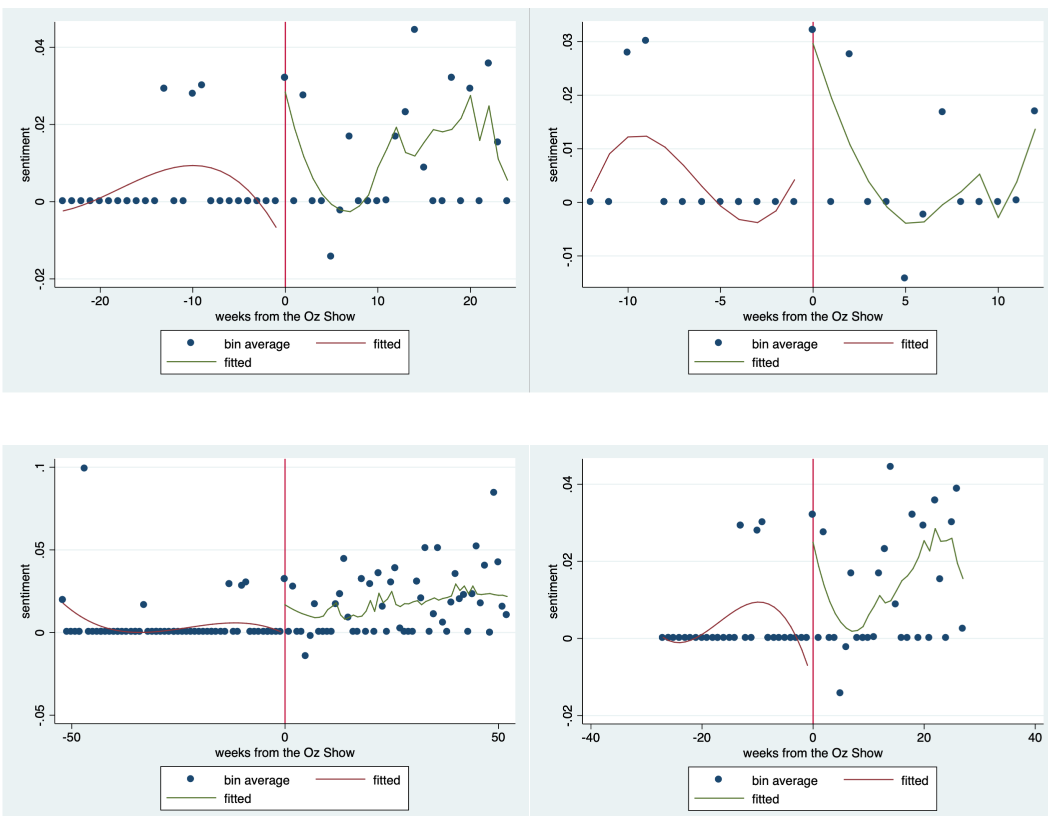

The findings shown in Figure 0.3 (for the running example) and column (1) of Table 2.12 imply that, the Oz Show leads to higher sentiment in the news articles. Though the language in news articles normally does not vary significantly on the emotion dimension, a higher sentiment and no correction indicate that news articles from legitimate outlets are not helping consumers in correcting the misleading healthcare information from the Oz Show. On the contrary, the increasing number of news articles carrying higher sentiment may even lead consumers to a false hope.

2) Consumer reviews

Now we examine the Oz effect on consumer reviews.

Bias correction

For bias correction, according to the column (2) of Table 2.11, we can see that the correction rate is also very low. However, compared to that of News articles, among all consumer reviews citing Dr. Oz or his show on Amazon, a larger proportion of UGC provides correction. For example, for 7-Keto, 0.0667/0.1=66.7% of reviews that cited Dr. Oz or his show contain correction or critiques. This indicates that many consumers are learning the lesson in an expensive way—trying the products. Despite the high conditional probability, the absolute correction rate is still very low: all less than 6.7%. This indicates that consumers may not be able to rely on the helpful UGC to get a more balanced view of the concerned product against the hype news.

Emotion

Regarding emotions, the results in column (1) of Table 2.13 and column (5) of Table 2.12 imply that the Oz Show leads to more negative emotion and less positive emotion in UGC. It means that consumers are expressing less excitement for the focal products and more disappointment or anger after experiencing the products, due to the high expectation of the weight-loss effect after getting the recommendation of Dr. Oz. The results with other bandwidths reported in the table show either the same findings or non-significant effects (but not the opposite significance). So the effect on emotion is not always robust but at least is not going the other direction.

Table 2.13. Oz Effect on Negative Emotion in UGC

	1 quarter bandwidth	2 quarter bandwidth	3 quarter bandwidth
After Oz	0.395***	-0.074	-0.038
	(4.77)	(-1.22)	(-0.68)
Ads.	-0.000	-0.000	0.000
	(-0.10)	(-0.22)	(0.07)
Price	-0.008***	-0.006***	-0.006***
	(-4.15)	(-4.63)	(-5.20)
Ranking	0.019	-0.002	-0.004
	(1.21)	(-0.13)	(-0.38)
Ingredient FE	Yes	Yes	Yes
Season FE	Yes	Yes	Yes
Observations	306	665	758
Adjusted R^2	0.34	0.101	0.104

Note: t-statistics in parentheses; all with 3rd order time polynomials; *** p<0.01, ** p<0.05, * p<0.1

<u>*Sentiment*</u>

The findings shown in column (2) of Table 2.12 indicate higher sentiment in UGC due to the Oz Show. Together with the findings about bias correction and emotion, we conclude that, although some consumers express their bad experience with the concerned products by writing reviews after the Oz Show, overall, UGC contains more supportive and positive information rather than correcting the bias from the hype news.

3) Research articles

Finally, we discuss the Oz effect on research articles.

Bias correction

Among all the research articles across all the ingredients, only one research article mentioned Dr. Oz and corrected the claim of Dr. Oz.[19] Therefore, for research articles, the correction rate is almost zero.

A study by the British Medical Journal (Korownyk et al. 2014), which was cited in the FTC hearing, points out that around half of Dr. Oz's suggestions are misleading or have no research foundation. However, this study was published after the hearing, around the end of 2014, beyond the time window of our study.

Sentiment and Emotion

For sentiment and emotion, the results in column (3) and column (6) of Table 2.12 indicate that the Oz Show does not significantly affect the sentiment and emotion contained in research articles. This is expected because research articles typically take a neutral and scientific way to illustrate their findings, so they rarely contain emotion.

2.5.1.3. Analysis of the Oz Effect with Synthetic Control

We also check the robustness of the previous findings using the Synthetic Control method. As explained ealier, to obtain a reasonably large control pool, good candidates of the treated categories include Safflower oil, Chitosan, Raspberry Ketones, 7-Keto, and Forskolin. For each of these five ingredients, we can form a control pool of at least four other ingredients (i.e., Green Coffee Bean, Moringa, Saffron Extract, and Garcinia Cambogia).

---

[19] Tessa Finney-Brown (2013), "Reviews of articles on medicinal herbs", Australian Journal of Herbal Medicine 2013 25(4).

Table 2.14. Comparison of Treated and Synthetic Control (for information intensity)

		Treated	Synthetic	Sample Mean
Safflower Oil	#news	0.846	0.115	0.058
	ads	0	0	10.445
Chitosan	#news	0.462	0.115	0.058
	ads	0	0	10.445
Raspberry Ketones	#news	0	0	0.087
	ads	0	41.707	10.445
7-Keto	#news	0.038	0.038	0.087
	ads	0	26.537	10.445
Forskolin	#news	0.038	0.038	0.087
	ads	0	24.296	10.445

The quantitative inference of the estimates is based on the placebo test akin to the classic permutation inference framework (Abadie et al. 2010). Table 2.14 compares the treated unit, the synthetic control unit, and the sample mean on news intensity and advertisement spending during the pre-intervention period.

Table 2.15 shows the results from synthetic control analysis for each of the five ingredients separatly.

The results from Synthetic Control analysis deliver consistent findings with those from the RDIT analysis. For four out of five categories, the Oz Effect on the number of news article as well as sentiment is significantly positive. The emotion effect is not significant at the 95% confidence level.

Table 2.15. Results Using Synthetic Control

	Dependent Var.	Estimate $\hat{\alpha}_{1t}$	95% CI
Safflower Oil	#news	1.1429***	0.6437, 1.6419
	Sentiment	0.0116**	0.0006, 0.0226
	Positive emotion	0.0255	-0.0090, 0.0600
Chitosan	#news	0.5000**	0.0064, 0.9935
	Sentiment	0.0116*	-0.0004, 0.0070
	Positive emotion	-0.0153*	-0.0092, 0.0602
Raspberry Ketones	#news	1.4439**	0.0005, 2.8873
	Sentiment	0.0197**	0.0021, 0.0373
	Positive emotion	0.0714*	-0.0062, 0.1491
Forskolin	#news	0.1692**	0.0248, 0.2887,
	Sentiment	0.0018**	0.0002, 0.0031,
	Positive emotion	-5.55E-06	-1.8344e-05, 7.2469e-06
7-Keto	#news	-0.1286	-0.3613, 0.1041
	Sentiment	-0.0009	-0.0021, 0.0003
	Positive emotion	-7.40E-06	-2.0370e-05, 5.5737e-06

Note: *** p<0.01, ** p<0.05, * p<0.1; Placebo tests are in appendix.

2.5.1.4. Possible Reasons for the above Findings

Our findings suggest that news articles are amplifying rather than correcting the misleading message originated from celebrity doctors. To rationalize these findings, we consider possible reasons from both the information supply side and the demand side.

Unlike fake news, the hype news generated by celebrity doctors is not easy to dis-prove. It requires research and raw data from scientific studies, which are not resources available to news agencies. As government officials noted in the FTC hearing,[20] "In response to our requests for scientific substantiation, companies usually will submit write-ups of human clinical studies,

[20] https://www.govinfo.gov/content/pkg/CHRG-113shrg92998/html/CHRG-113shrg92998.htm

sometimes published in peer-reviewed journals. While these studies may appear facially plausible, in a number of cases, we have discovered serious flaws, or worse, outright fabrications once we obtain the underlying data." However, media companies are not able to get these raw data to make in-depth investigation. Moreover, related to the media bias literature (Gentzkow& Shapiro 2006; Gentzkow & Shapiro 2010), media companies care about profitibility. Facing a high cost of conducting independent investigation versus a high opportunity to attract audience using the hype information, most of them will choose the latter. This points out the importance of government intervention.

On the demand side of public health information, according to Gentzkow and Shapiro (2006, 2010), consumers have prior beliefs and subjective ideology. People who are longing for weight loss remedies are more inclined to believe the bright side of a product rather than the downside. To cater to this consumer preference, news articles may choose to contain more positive sentiment in the content, rather than try to wake up consumers with false hope.

Unfortunately, we do not have data to test which reason is more plausible. We leave this as an open question for future research.

2.5.2. Market response to hype news

The previous findings speak to the Oz Effect on the supply of public information for the concerned products. A natural follow-up question is how the consumers react to the hype news and whether they demand more information after the Oz Show. Though we cannot directly demonstrate the linkage between the information supply and demand, we can infer this link by examining how consumer search is affected by the Oz Show. The other question worth asking is how firms react to the hype news. That is, whether there is increasing product supply or any price change caused by the hype news. Answering this question helps us get a sense of how hype news would impact consumers in the market if they decided to purchase the focal products as a result of following the biased information. We answer these two questions in this section.

Figure 0.4. Market Response to Hype News

Garcinia Cambogia (Nov, 2012)	Green Coffee Bean (April, 2012 and Sep, 2012)	Raspberry Ketone (Feb, 2012)

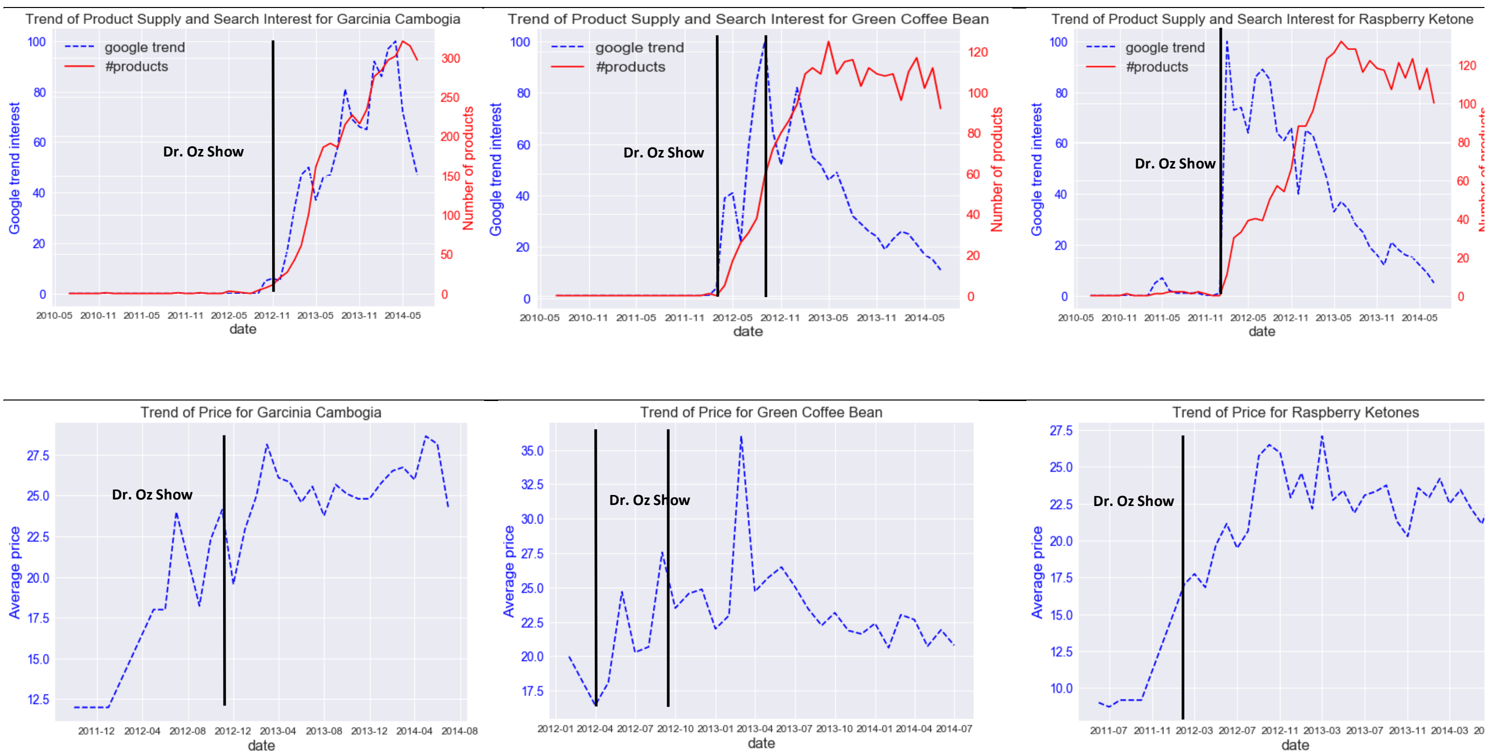

Figure 0.4 shows the trends of product supply, consumer search, and price changes for three representative ingredients mentioned on the Oz Show, including Garcinia Cambogia, Green Coffee Bean, and Raspberry Ketones. The other ingredients also share similar patterns for the

corresponding trends. Below we discuss the implication of each trend and conduct RDIT analysis to quantify the corresponding Oz Effect.

Consumer search interest

The dotted curve in the first row of Figure 0.4 are Google Trend for each ingredient, describing the intensity of consumer search of the focal key word over time. The verticle line points to the air date of the Oz Show featuring the focal ingredient. It is clear that consumers got interested in knowing more about the ingredient thus searched for the related information after the show. This increase in consumer search is a direct demand for information, which could also convert to demand for products that contain the concerned ingredient.

The figure also indicates that there is some time lag between searching (Google trend) and purchasing (product supply). Though both of them increase after *The Dr. Oz Show*, the increase of searching interest happens more promptly. This evidence in Figure 0.4 indicates the possibility that consumers' purchase decisions may be affected by the information from other sources after *The Dr. Oz Show* rather than the show alone.

Consistently, the RDIT analysis conducted at monthly level indicates that the Oz Show induced consumer search interest about the concerned product.[21] In the third column of Table 2.16, we can see that the change in magnitude is huge: in the scale from 0 to 100, the Oz Show increased search volumn by 27.4, which is a big surge in consumer demand for information.

Supply change

[21] It is studies at a monthly level because the Google search trend data granularity we obtained is at month level.

From the change in number of products shown in the first row of Figure 0.4, we can see that the increase in the supply of products is phenomenal, from almost around zero to several hundred listings on Amazon.com. However, the RDIT analysis in Table 2.16-column (1) indicates that The Dr. Oz Show is not causing any immediate supply-side response. Firms may not base their product strategy on a single show, instead, they move only when they see clear consumer needs. So, as the running example shown in Figure 0.5, we see the product offerings increase gradually over time, not abruptly after the Oz Show.

Figure 0.5. Oz Effect on Product Supply (e.g., Green Coffee Bean)

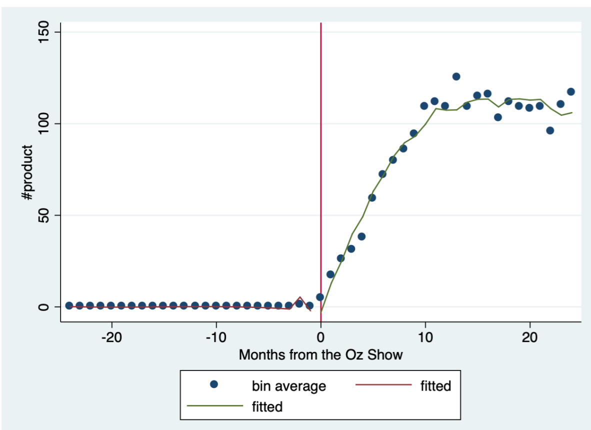

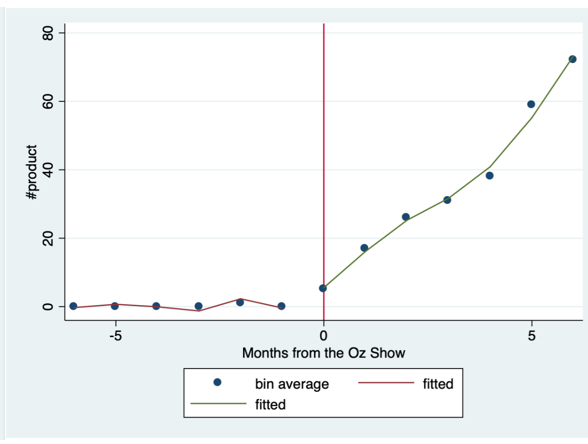

<u>Price response</u>

The second row of Figure 0.4 presents the price trend. We can see that, for some ingredients, the price seems to be higher on average after the Oz Show. We also conduct RDIT analysis with price as the outcome variable. According to the results in the second column of Table 2.16, the Oz Show leads to significantly higher price for the concerned products. This result implies that even though product offering seems not to respond to the Oz Show immediately, due to possible reasons like production delay or budget constraint, firms forsee a increase in consumer demand and respond by increasing the price of the current products, which can be implemented with little cost and delay.

Table 2.16. Oz Effect on Product Offering, Price, and Consumer Search

	(1) #Product	(2) Price	(3) Consumer Search
After Oz	4.434	5.658**	27.424***
	(0.33)	(2.21)	(4.16)
Ads.	0.052	0.028	-0.003
	(0.33)	(0.92)	(-0.15)
Price	1.165***		
	(5.00)		
#Product		0.044***	
		(5.00)	
Ingredient FE	Yes	Yes	Yes
Season FE	Yes	Yes	Yes
Observations	486	486	421
Adjusted R^2	0.492	0.579	0.448

Note: t-statistics in parentheses; Two years before and after the Oz Show; all with 3rd order time polynomials; *** p<0.01, ** p<0.05, * p<0.1

2.6. Conclusion

This paper aims to understand how misleading health information source affects subsequent information generation. We combine causal inference and natural language processing to study how biased information from celebrity doctors affect the supply of other publicly available health information for over-the-counter (OTC) weight loss products. Our textual analysis employs both traditional machine learning and state-of-the-art deep learning natural language processing techniques. We use regression discountinuity in time design and synthetic control to uncover the Oz Effect on information intensity and content (captured by sentiment, positive emotion, negative emotion, and bias correction).

Our findings indicate that the Oz Show caused increasing news coverage from legitimate outlets. Surprisingly, we find that the news articles generated after the Oz Show are written with higher sentiment, no significant change in emotions, and little correction for the ingredients peddled by the celebrity doctor, meaning that news articles are acting as a magnifier rather than rectifier. While the government and firms like Google, Baidu, and Facebook, have been fighting against fake news in the healthcare domain, our finding suggests that even legitimate news articles respond to the biased healthcare information by propagating and magnifying it, rather than correcting it. Regarding the more credible information source, namely scholarly journals, only one out of thousands of peer-reviewed articles directly correct what Dr. Oz said on the talk show. For user-generated content, though some consumers reviews provides correction and shows increasing negative emotion, it is overwhelmed by the larger amount of UGC that supports the misleading information. Moreover, we find that the hype news from The Dr. Oz Show leads to about 30% more consumer search for the recommended ingredients. The price of the featured products also significanlty increase due to the recommendation on the Oz Show. The findings all together tell us that consumers try to look for more information after listening to hype news, but what they can get from publicly available sources are predominately supportive of the biased information. As a result, we may face the unfortunate situation where consumers suffer from both health damage and financial cost.

This paper makes a substantive contribution to provide concrete empirical evidence on how biased and hyped health information source affects subsequent information generation from other sources. The results have important public policy and managerial implications. For public policymakers, our study sheds light on how to supervise health-related media content, in order to protect consumers from misleading healthcare information and help their decision-making. Managerially, our finding sheds light for how to manage new product provision according to public information in the OTC market.

There are a few limitations and directions for future research. First, our research is most relevant to over-the-counter healthcare products. The results may not apply to prescription-based medicine. Given patients will access very different information set in these two scenarios, the effect of biased messages sent by celebrity doctors on subsequent market response may be quite different. Second, the tone of the Oz Show may change over time, therefore, the effect can vary over time and across ingredients. As we restrict the ingredients in our study to those mentioned on the show within one year (2012), the episodes featuring these ingredients share a high degree of similarity, so the variation in the tone may be less of a concern. Lastly, due to the data limitation, we do not observe the sales for each product thus cannot quantify the Oz Effect on sales. This can be a future research direction, if new data can provide a precise measure of sales.

Reference

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. 2015. "Comparative Politics and the Synthetic Control Method." American Journal of Political Science 59 (2): 495–510.

Allcott, Hunt, and Matthew Gentzkow. 2017. "Social Media and Fake News in the 2016 Election." Journal of Economic Perspectives 31 (2): 211–36.

Autor, David, Gary Becker, Gary Chamberlain, Raj Chetty, Tim Conley, Liran Einav, Edward Glaeser, et al. 2010. "What Drives Media Slant? Evidence from u.s. Daily Newspapers." Econometrica 78 (1): 35–71.

Badie, Alberto A, Alexis D Iamond, and Jens H Ainmueller. 2010. "Synthetic Control Methods for Comparative Case Studies : Estimating the Effect of California's Tobacco Control Program." Journal of the American Statistical Association 105 (490): 493–505.

Bates, B. R., Romina, S., Ahmed, R., & Hopson, D. (2006). The effect of source credibility on consumers' perceptions of the quality of health information on the Internet. Medical Informatics and the Internet in Medicine, 31(1), 45–52.

Blei, D. M., Edu, B. B., Ng, A. Y., Edu, A. S., Jordan, M. I., & Edu, J. B. (2003). Latent Dirichlet Allocation. *Journal of Machine Learning Research*, *3*, 993–1022.

Ching, A. T., Clark, R., Horstmann, I., & Lim, H. (2015). The Effects of Publicity on Demand: The Case of Anti-Cholesterol Drugs. *Marketing Science*, (April), 150817104945009.

Chintagunta, P. K., Jiang, R., & Jin, G. Z. (2009). *Information, learning, and drug diffusion: The case of Cox-2 inhibitors*. *Quantitative Marketing and Economics* (Vol. 7).

D'Andrea, A., Ferri, F., Grifoni, P., & Guzzo, T. (2015). Approaches, tools and applications for sentiment analysis implementation. *International Journal of Computer Applications*, *125*(3), 26–33.

Dahlberg, L. (2001). Computer-Mediated communication and the public sphere: A critical analysis. *Journal of Computer-Mediated Communication*, *7*(1), 0.

Eisenberg, M. (2008). The Impact of Over the Counter (OTC) Weight Loss Product Television Advertisements on the Consumption of OTC Weight Loss Products and Other Diet-Related Behaviors. *Dissertation*.

Eysenbach, G., & Köhler, C. (2002). How do consumers search for and appraise health information on the world wide web? Qualitative study using focus groups, usability tests, and in-depth interviews. *British Medical Journal*, *324*(7337), 573.

Eysenbach, G. (2007). *Credibility of Health Information and Digital Media: New Perspectives and Implications for Youth*. *The John D. and Catherine T. MacArthur Foundation Series on Digital Media and Learning*.

Friggeri, Adrien, Dean Eckles, and Justin Cheng. 2014. "Rumor Cascades." In Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media Adrien.

Gentzkow, Matthew, Chicago Booth, Jesse M Shapiro, Chicago Booth, and Daniel F Stone. 2014. “Media Bias in the Marketplace : Theory.” https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2393593.

Gentzkow, Matthew, and Jesse M Shapiro. 2006. “Media Bias and Reputation.” Journal of Political Economy 114 (2).

Ghose, A., & Ipeirotis, P. G. (2011). Estimating the Helpfulness and Economic Impact of Product Reviews. *International Journal of Innovative Research and Development*, *23*(10), 1498–1512.

Goel, S., Watts, D., & Goldstein, D. (2012). The Structure of Online Diffusion Networks. *Proceedings of the 13th ACM Conference on Electronic Commerce*, 623-638.

Gong, J., Abhishek, V., & Li, B. (2017). Examining the Impact of Contextual Ambiguity on Search Advertising Keyword Performance : A Topic Model Approach. *MIS Quarterly*, (July), 1–68.

Hausman, Catherine, and David S Rapson. 2017. “Regression Discontinuity in Time : Considerations for Empirical Applications.”

Hu, Y., & Sundar , S. S. (2010). Effects of Online Health Sources on Credibility and Behavioral Intentions. *Communication Research*, *37*(1), 105–132.

Kalra, a., Li, S., & Zhang, W. (2011). Understanding Responses to Contradictory Information About Products. *Marketing Science*, *30*(December 2013), 1098–1114.

Khan, R., Misra, K., & Singh, V. (2016). Will a Fat Tax Work? *Marketing Science*, *35*(1), 10–26.

Kim, Y. (2014). Convolutional Neural Networks for Sentence Classification. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*.

Korownyk, Christina, Michael R. Kolber, James McCormack, Vanessa Lam, Kate Overbo, Candra Cotton, Caitlin Finley, et al. 2014. "Televised Medical Talk Shows - What They Recommend and the Evidence to Support Their Recommendations: A Prospective Observational Study." *British Medical Journal* 349 (December): 1–9. doi:10.1136/bmj.g7346.

Lappas, T., Sabnis, G., & Valkanas, G. (2016). The Impact of Fake Reviews on Online Visibility : A Vulnerability Assessment of the Hotel Industry. *Information Systems Research*, *1*(December), 22.

LeCun, Y., Bengio, Y. & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.

Lee, D., Hosanagar, K., & Nair, H. S. (2017). Advertising Content and Consumer Engagement on Social Media: Evidence from Facebook. Forthcoming, *Management Science*.

Lee, T. Y., & BradLow, E. T. (2011). Automated Marketing Research Using Online Customer Reviews. *Journal of Marketing Research*, *48*(5), 881–894.

Lee, David S, and Thomas Lemieux. 2010. “Regression Discontinuity Designs in Economics.” Journal of Economic Literature 48 (June): 281–355.

Liu, X., Lee, D., & Srinivasan, K. (2017). Large Scale Cross Category Analysis of Consumer Review Content on Sales Conversion Leveraging Deep Learning. *Working Paper*, 1–48.

Metzger, M. J., & Flanagin, A. J. (2011). Using Web 2.0 Technologies to Enhance Evidence - Based Medical Information. *Journal of Health Communication: International Perspectives*, *16*(1), 45–58.

Porter, L. and Golan, G. J. (2006). From Subservient Chickens to Brawny Men: a Comparison of Viral Advertising to Television Advertising. *Journal of Interactive Advertising,* 6 (2).

Resnik, A., & Stern, B. L. (1977). An Analysis of Information Content in Television Advertising. *Source Journal of Marketing*, *41*(1), 50–53.

Senecal, S., & Nantel, J. (2004). The influence of online product recommendations on consumers’ online choices. *Journal of Retailing*, *80*, 159–169.

Teisl, M. F., Bockstael, N. E., & Levy, A. (2001). Measuring the Welfare Effects of Nutrition Information. *American Journal of Agricultural Economics*, *83*(1), 133–149.

Timoshenko, A., & Hauser, J. R. (2017). Identifying Customer Needs from User-Generated Content. *Working Paper*.

Tirunillai, S. & Tellis , G. J. (2012). Does chatter really matter? Dynamics of user-generated content and stock performance. *Marketing Science,* 31(2), 198–215.

Tirunillai, Seshadri, and Gerard J Tellis. (2017). “Does Offline TV Advertising Affect Online Chatter? Quasi-Experimental Analysis Using Synthetic Control.” Marketing Science 36 (6): 862–78.

Uetake, K., & Yang, N. (2017). Success Breeds Success : Weight Loss Dynamics in the Presence of Short-Term and Long-Term Goals. *Working Paper*.

Vasterman, P. L. M. (2005). Media-Hype. *European Journal of Communication*, 20(4), 508–530.

Vosoughi, Soroush, Deb Roy, and Sinan Aral. (2018). “The Spread of True and False News Online.” Science 359: 1146–51.

Wang, P., Xu, J., Xu, B., Liu, C., Zhang, H., Wang, F., & Hao, H. (2015). Semantic Clustering and Convolutional Neural Network for Short Text Categorization. *Proceedings ACL 2015*, 352–357.

Ye, Q., Law, R., & Gu, B. (2009). International Journal of Hospitality Management The impact of online user reviews on hotel room sales. *International Journal of Hospitality Management*, *28*, 180–182.

Yildirim, P., Gal-Or, E., & Geylani, T. (2013). User-Generated Content and Bias in News Media. *Management Science*, *1909*(January 2015), 1–12.

Zhang, Y., & Wallace, B. (2016). A Sensitivity Analysis of (and Practitioners' Guide to) Convolutional Neural Networks for Sentence Classification. *arXiv Preprint*, arXiv:1510.03820v4.

Zuckerman, D. (2003). Hype in health reporting: "checkbook science" buys distortion of medical news. *International Journal of Health Services*, *33*(2), 383–389.

Appendix

2.A. Survey Form Designed in Amazon Mechanical Turk

We are studying textual information about weight-loss products.

You will see a text. The following questions ask if a certain type of content is present in the review.

Please read the text and questions carefully before answering the questions.

Please answer the following questions about the review.

1. Have you heard of Dr. Oz?

Yes ◯ No ◯

2. Have you ever watched the Dr. Oz Show?

Yes ◯ No ◯

3. According to this customer review, does this person like or dislike Dr. Oz or the Dr. Oz Show?

Like ◯ Dislike ◯ Neutral ◯

4. According to this customer review, does this person believe what Dr. Oz said is correct?

Believe ◯ Not believe ◯ Neutral ◯

8. ***[Emotion]*** Does this text include emotional content? (e.g., "I feel so sad that it has no effect on me.")

Yes ◯ No ◯

If you answered yes above, judge if it is a positive or negative emotion. If answered no above, then select Not Applicable.

Positive emotion ◯ Negative emotion ◯ Not Applicable ◯

2.B. Journal Names

Source of Research Articles: Peer Reviewed Journals

JOURNAL NAME	JOURNAL NAME	JOURNAL NAME	JOURNAL NAME
Journal of Materials Science : Materials in Medicine	Emirates Journal of Food and Agriculture	Clean Technologies and Environmental Policy	Australian Journal of Herbal Medicine
PLoS One	The Journal of Dairy Research	Journal of Biological Engineering	Biodegradation
International Journal of Molecular Sciences	Critical Reviews in Food Science and Nutrition	European Food Research and Technolog	Anti - Corrosion Methods and Materials
Marine Drugs	Gene Therapy	Virology Journal	Molecular Biology Reports
Journal of Nanomaterials	Indian Journal of Pharmaceutical Sciences	Polymer Composites	Antonie van Leeuwenhoek
Journal of Food Science and Technology	Tropical Animal Health and Production	Journal of Bone and Mineral Metabolism	Oxidative Medicine and Cellular Longevity
Lipids	Carcinogenesis	Nutrition Research Reviews	PPAR Research
International Journal of Polymer Science	Journal of Industrial Microbiology & Biotechnology	Current Topics in Nutraceuticals Researc	AIDS Research and Therapy
Journal of Polymers and the Environment	Obesity	PLoS Neglected Tropical Diseases	Pakistan Journal of Zoology
Pharmaceutical Research	The Proceedings of the Nutrition Society	Public Health Nutrition	Al Ameen Journal of Medical Sciences
The British Journal of Nutrition	International Journal of Environmental Science and Technology	International Journal of Obesity	Biomaterials Research
Journal of Dairy Science	International Journal of Food Science and Technology	Polymer Engineering and Science	Notulae Scientia Biologicae
Journal of Polymer Materials	Journal of Food Protection	Kidney International	Biomedical Microdevices
Scientific Reports (Nature Publisher Group)	Journal of Food Science	Diabetology & Metabolic Syndrome	Advances in Materials Science and Engineering
BioMed Research International	Journal of Dentistry	Diabetes	Biophysics
African Journal of Biotechnology	Tissue Engineering	Water Environment Research	Acta Agron—mica
Vaccine	International Journal of Biomaterials	Agroforestry Systems	Acta Chimica Slovaca
Nanomedicine	Molecular Therapy	Journal of Nanobiotechnology	BioResearch Open Access

Source of Research Articles: Peer Reviewed Journals (Continued)

JOURNAL NAME	JOURNAL NAME	JOURNAL NAME	JOURNAL NAME
Nanomedicine	Molecular Therapy	Journal of Nanobiotechnology	BioResearch Open Access
Environmental Science and Pollution Research International	Chemistry Central Journal	Analytical and Bioanalytical Chemistry	Journal of the International Society of Sports Nutrition
Journal of Animal Science	Journal of Young Pharmacists	Clinical Lipidology	BMC Plant Biology
Applied Mechanics and Materials	World Journal of Life Sciences and Medical Research	Pigment & Resin Technology	BMC Genomics
International Journal of Pharmaceutical Sciences and Research	Environmental Chemistry Letters	Journal of Nanotechnology	The Journal of Nutrition
Bioprocess and Biosystems Engineering	Biotechnology Letters	Alternative Therapies in Health and Medicine	Journal of Strength and Conditioning Research
Nutrition	International Journal of Plastics Technology	Microbial Cell Factories	Ethiopian Journal of Environmental Studies and Management
Obesity Research	Caries Research	PLoS Pathogens	Journal of Nutrition and Metabolism
BMC Complementary and Alternative Medicine	Cancer Letters	Pakistan Journal of Medical Sciences Quarterly	Gastroenterology Research and Practice
International Food Research Journal	Acta Pharmaceutica	Research in Veterinary Science	Gesunde Pflanzen
Applied Microbiology and Biotechnology	Annals of Biomedical Engineering	Nutrition Reviews	Journal of Applied Poultry Research
European Food Research and Technology	Archives of Virology	Nutrition Journal	Euphytica
Applied Biochemistry and Biotechnology	European Journal of Plant Pathology	Nutrition and Cancer	Archives of Pharmacy Practice
European Journal of Nutrition	Journal of Biomedicine and Biotechnology	Nutrition & Metabolism	Amino Acids
Applied Biochemistry and Microbiology	Materials Science and Technology	Research Journal of Pharmacy and Technology	Annals of Nutrition & Metabolism
Indian Journal of Clinical Biochemistry	Journal of Conservative Dentistry	Nephron	International Journal of Cultural Property
Biotechnology and Bioprocess Engineering : BBE	International Journal of Carbohydrate Chemistry	Mycorrhiza	Agricultural History

Source of Research Articles: Peer Reviewed Journals (Continued)

JOURNAL NAME	JOURNAL NAME	JOURNAL NAME	JOURNAL NAME
In Vitro Cellular & Developmental Biology	Journal of Membrane Biology	Polish Journal of Veterinary Sciences	Journal of Insect Behavior
Global Journal of Research on Medicinal Plants	Journal of International Dental and Medical Research	Journal of Applied Microbiology	Age
Lipids in Health and Disease	Asian Journal of Research in Chemistry	Metabolic Brain Disease	International Journal of Food Sciences and Nutrition
Molecular Medicine Reports	Bulletin of Experimental Biology and Medicine	Journal of Sports Medicine and Physical Fitness	International Journal of Molecular Medicine
Iranian Journal of Basic Medical Sciences	Biological Trace Element Research	The Veterinary Record	Maejo International Journal of Science and Technology
The Pharma Innovation	Journal of Electronic Materials	Journal of Mammary Gland Biology and Neoplasia	Human and Experimental Toxicology
International Journal of Obesity and Related Disorders	Evidence-Based Complementary and Alternative Medicine	Science International	Journal of Toxicology
Asian Journal of Pharmaceutics	Experimental and Therapeutic Medicine	Journal of Diabetes & Metabolic Disorders	Italian Journal of Food Science
European Journal of Clinical Nutrition	Journal of Basic and Clinical Physiology and Pharmacology	Journal of Clinical Pharmacy and Therapeutics	Journal of Pharmaceutical Sciences and Research
Animal: an International Journal of Animal Bioscience	Journal of Pharmaceutical Education and Research	Journal of Bioenergetics and Biomembranes	Journal of Ayurveda and Integrative Medicine
Plant Growth Regulation	Journal of Materials Research	Focus On Geography	Journal of Biological Research
The American Journal of Surgery	International Journal of Photoenergy	Journal of Medical Toxicology	Theatre Research International
The Scientific World Journal	Metabolomics	Poultry Science	The Protein Journal
Food Biophysics	Circulation	BMJ Open	Journal of Mountain Science
Polymers & Polymer Composites	Pharmacognosy Communications	Asia Pacific Journal of Clinical Nutrition	Annual Review of Nutrition
Malaysian Journal of Pharmaceutical Sciences	Journal of Ocean University of China. JOUC	American Journal of Rhinology & Allergy	Journal of the American Dietetic Association
International Journal of Applied Science and Engineering	Fish Physiology and Biochemistry	Current Topics in Nutraceuticals Research	Breast Cancer Research and Treatment
Diabetes Care	Glycoconjugate Journal	Defence Science Journal	Diabetologia
Indian Journal of Medical Research	The Journal of Agricultural Science	International Journal of Aquaculture	International Archives of Allergy and Immunology

2.C. Robustness Check with Synthetic Control

The table below reports the optimal weights allocated to each units in the control pool, for each treated unit.

Table 2C.17. Optimal Weights Allocated to Units in Control Pool

	Treated					Control Pool
	Weights					
	Safflower Oil	Chitosan	RaKe	7-Keto	Forskolin	
#news	1	1	0.001	0.172	0.168	GaCa
	0	0	0	0.077	0.082	GreCof
	0	0	0.001	0.116	0.168	moringa
	0	0	0.998	0.635	0.582	saffr
Sentiment	1	1	0.001	0.202	0.208	GaCa
	0	0	0	0.134	0.155	GreCof
	0	0	0.001	0.190	0.235	moringa
	0	0	0.998	0.474	0.402	saffr
Positive Emotion	0	0	0.997	0.997	0.997	GaCa
	0.5	0.5	0	0	0	GreCof
	0.5	0.5	0	0	0	moringa
	0	0	0.003	0.003	0.003	saffr
Negative Emotion	\					

Note: "\" means there is no variation in the variable, thus no results are obtained.

The main text analyzes five treated ingredients separately, as they have reseasonably many control candidates in the pool. If we want to enlarge the control pool by including all the ingredients, without considering whether it has been mentioned by the Oz Show or not, then we can examine more units. Below is the result for Garcinia Cambogia.

Figure 2.C1. Effect on Information Intensity

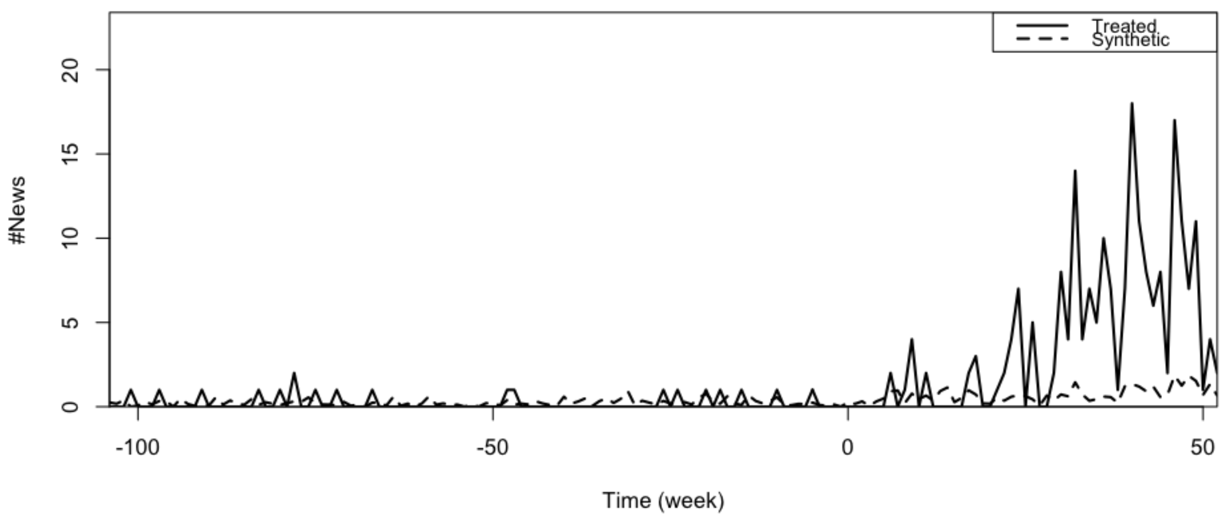

The above figure shows the trend of the number of news articles generated every week for both the treated and the synthetic unit (e.g., Garcinia Cambogia). Time 0 is when the treated unit (an ingredient) was recommended on the Dr. Oz Show.

We choose the number of news articles generated at weekly level in the past year as our predictor variables for constructing the synthetic control. For the case of Garcinia Cambogia, we can get the following weights

weights	Unit names
0.07	chitosan
0.056	cla
0.105	for
0.07	glu
0.046	GreCof
0.131	keto
0.105	moringa
0.044	RaKe
0.105	saffr
0.068	safl
0.2	yacon

2.D. RDIT Results with Various Bandwidths

Table 2.D1. Oz Effect on Sentiment of News Articles (Pool)

	24 periods window	12 periods window	52 periods window	36 periods window
VARIABLES	avg_sentiment	avg_sentiment	avg_sentiment	avg_sentiment
After Oz	0.009**	0.008	0.003	0.008**
	(2.20)	(1.16)	(0.74)	(2.10)
total_ads	0.000**	0	0.000***	0.000***
	-2.53	(-0.36)	-4.71	-4.37
Ingredient FE	Yes	Yes	Yes	Yes
Season FE	Yes	Yes	Yes	Yes
Observations	595	306	1,264	669
Adjusted R^2	0.123	0.088	0.226	0.134

Note: t-statistics in parentheses; all with 3rd order time polynomials; *** p<0.01, ** p<0.05, * p<0.1

Table 2.D2. Oz Effect on Sentiment in UGC

	24 periods window	12 periods window	52 periods window	36 periods window
VARIABLES	avg_sentiment	avg_sentiment	avg_sentiment	avg_sentiment
After Oz	0.131***	0.023	0.089***	0.125***
	(2.72)	(0.27)	(2.98)	(2.83)
total_ads	-0.000	-0.000	-0.000	-0.000
	(-1.20)	(-0.69)	(-0.96)	(-1.35)
avg_price_week	0.003**	0.005**	0.003***	0.002**
	(2.37)	(2.41)	(4.08)	(2.16)
avg_rank	0.015	-0.006	0.008	0.014
	(1.45)	(-0.36)	(1.26)	(1.56)
Ingredient FE	Yes	Yes	Yes	Yes
Season FE	Yes	Yes	Yes	Yes
Observations	665	306	1,583	758
Adjusted R^2	0.104	0.082	0.065	0.097

Note: t-statistics in parentheses; all with 3rd order time polynomials; *** p<0.01, ** p<0.05, * p<0.1

Table 2.D3. Oz Effect on Sentiment in Research Articles

VARIABLES	24 periods window avg_sentiment	12 periods window avg_sentiment	6 periods window avg_sentiment
After Oz	-0.004	-0.029	0.025
	(-0.20)	(-1.01)	(0.48)
total_ads	0.000***	0.000***	0.000
	(4.80)	(5.07)	(1.59)
Ingredient FE	Yes	Yes	Yes
Season FE	Yes	Yes	Yes
Observations	488	250	130
Adjusted R^2	0.231	0.275	0.391

Note: t-statistics in parentheses; all with 3rd order time polynomials; *** p<0.01, ** p<0.05, * p<0.1

Table 2.D4. Oz Effect on Positive Emotion in UGC

VARIABLES	24 periods window avg_emotionP	12 periods window avg_emotionP	36 periods window avg_emotionP
After Oz	-0.215***	0.128	-0.264***
	(-2.83)	(1.29)	(-3.78)
total_ads	0.000	-0.000	0.000
	(0.84)	(-0.64)	(0.81)
avg_price_week	0.004**	0.006**	0.002
	(2.47)	(2.54)	(1.50)
avg_rank	0.037**	0.003	0.035**
	(2.32)	(0.16)	(2.39)
Ingredient FE	Yes	Yes	Yes
Season FE	Yes	Yes	Yes
Observations	665	306	758
Adjusted R^2	0.236	0.268	0.237

Note: t-statistics in parentheses; all with 3rd order time polynomials; *** p<0.01, ** p<0.05, * p<0.1

Table 2.D5. Oz Effect on Positive Emotion in Research Articles

VARIABLES	24 periods window avg_emotionP	12 periods window avg_emotionP	6 periods window avg_emotionP
After Oz	0.058*	-0.013	0.041
	(1.81)	(-0.27)	(0.55)
total_ads	0.000	0.000	0.000
	(0.36)	(0.82)	(1.18)
Ingredient FE	Yes	Yes	Yes
Season FE	Yes	Yes	Yes
Observations	488	250	130
Adjusted R-squared	0.041	0.030	0.094

Note: t-statistics in parentheses; all with 3rd order time polynomials; *** p<0.01, ** p<0.05, * p<0.1

Chapter 3

Does Fast Fashion Increase the Demand for Premium Brands? A Structural Analysis

3.1. Introduction

The global fashion industry has reached an estimated value of 3 trillion dollars.[22] Traditional luxury fashion brands, such as Gucci, Prada, and Louis Vuitton, have maintained a strong position within the industry, backstopped by the increasing demand from developing economies such as China. At the same time, fast fashion brands such as Zara, Forever 21, and H&M have been storming the globe with their versatile styles and low price. The path to success of these fast fashion brands, however, is nothing short of controversy. Every year, large fast fashion chains spew close-to-the-runway originals at lightning speed. On the one hand, the high-end brands, believing these copycats will steal their customers and hurt their profitability, spare no effort in fighting back by launching lawsuits against them.[23] On the other hand, high-end brands may not face any threat if their consumers and those of fast fashion brands are different segments with variant values of brands and styles. In addition, lawmakers tend to view the utilitarian nature of clothing and fashion as more important than its artistic and stylistic purposes; therefore fashion

---

[22] https://fashionunited.com/global-fashion-industry-statistics

[23] For example, https://wwd.com/business-news/legal/the-5-five-biggest-lawsuits-facing-fashion-retail-10875211/

designs are not under the protection of copyright law.[24] In spite of this tension on the enforcement and effects of copycats, the effect of fashion copycats on high-end brands remains empirically unclear.

In this paper, we estimate the impact of low-end copycats on the demand for high-end brands. We develop a dynamic structural model of individual consumer's fashion choices, which allows for counterfactual analysis of alternative copyright policies against copycats. Contrary to the conventional wisdom, we find that prohibiting low-end copycats can decrease the demand of high-end brands significantly. We find novel mechanisms contributing to this result, which are distinct from the promotional effect documented in the counterfeit literature (e.g., Qian, 2014): first, fewer style choices from low-end brands would limit the mix-and-match choices for consumers and put them on greater financial constraint to get a satisfactory ensemble of clothes, resulting in them buying less high-end brands; second, the lack of good styles from low-end brands will make it harder for consumers to build up their popularity/likeability, which limits the complementary value of high-end brands. As a result, consumers adopt less high-end brands. Our findings suggest that the above-mentioned market expansion effect dominates the competition effect. Other counterfactual analyses examine the consequence if the brand or peer feedback cannot be seen, which is the case for many other social medium and offline markets.

We overcome substantial technical and empirical challenges to obtain these results. Traditionally, there are two challenges to studying the micro-level consumer choices of fashion goods. First, fashion styles are not quantifiable. Second, individual-level choices on fashion

---

[24] http://www.thefashionlaw.com/home/how-do-fast-fashion-retailers-get-away-copying-high-fashion-brands

brands and styles across a large pool of brands are not available. In this paper, we employ state-of-the-art deep learning techniques to quantify fashion styles from fashion images. We overcome the second challenge by studying the choices over brands and styles for fashion conscious consumers on social media. Nowadays, fashion is one of the most popular contents generated by users on social media (Hu et al. 2014). More and more fashion consumers post what they wear online, and importantly, these social media users become trendsetters and influence a large number of other fashion consumers.[25] Therefore, investigating how these consumers make choices on brand and style can help us understand the market demand of fashion goods. Our data is from a large online fashion-sharing platform where users post their fashion pictures and evaluate others' pictures. The data comprise 10262 active users and 64681 fashion posts and span over three years. We account for consumer heterogeneity and estimate the structural model following a Hierarchical Bayesian framework.

Substantively, our results have implementable policy implications to both managers as well as policymakers. Managerially, we provide novel insights on how copycats can help the high-end brands, which guide their product strategy. In fact, some high-end brands have started to produce their own low-end frugal version of similar styles, consistent with the first mechanism of copycat effects in our findings. Moreover, for fashion companies, understanding how fashion consumers value brands and styles can help managers infer the market demand and make the optimal investments in branding and product design. From the policy-making perspective, we provide novel insights on the potential consequence of alternative copyright policies for fashion designs.

---

[25] http://www.latimes.com/fashion/la-ig-bloggers-20160809-snap-story.html

More generally, our findings speak to the debate on whether copyright or patent protection encourages or discourages innovation in the fashion industry: with more demand brought by low-end copycats, companies can get more money to invest in innovation, which may lead to more creative designs for the entire fashion market.

This paper also contributes methodologically and theoretically. Methodologically, we make two contributions. First, we develop a framework to analyze consumer choices where visual features are important product attributes and other people's opinions heavily affect the decision-making. Second, we use deep learning and image processing techniques to quantify fashion styles to make the analysis of fashion style choices possible. Theoretically, our findings provide new insights on how copycat products can benefit the original brands, which also apply to the cases of counterfeits and pirated goods if consumers have mix-and-match choices and popularity concern.

The rest of the paper is organized as follows. In 3.2, we review the literature related to this paper. Section 3.3 presents the raw data, visual feature extraction, and exploratory analysis. Section 3.4 describes our model. We illustrate the identification and estimation strategy in Section 3.5. We report the estimated results in Section 3.6, followed by the counterfactual analysis in Section 3.7. Section 3.8 concludes.

3.2. Literature Review

Our study relates to marketing and economics literature on branding, counterfeits and piracy, conspicuous good consumption, as well as the literature on machine learning methods and applications.

As we seek to examine the brand value and style value for fashion goods, our paper is related to the marketing literature on branding (e.g., Borkovsky et al., 2017; Goldfard, 2009; Keller & Lehmann, 2006; Kamakura & Russell, 1993). More recently with unstructured data, Nam et al. (2017) investigate the qualitative brand information harvested from social tags in the textual form. Liu et al. (2018) study consumers' brand perception on social media using visual data. In this paper, we focus on fashion goods, specifically clothing. We examine how consumers value brands versus styles, and how copycat styles affect the demand of premium brands measured by the units of clothing items adopted in the social media posts.

This paper relates to the literature of counterfeits and piracy (e.g., Qian, 2014; Ma et al., 2016; Oberholzer-Gee & Strumpf, 2007; Smith & Telang, 2009), which provided evidences of both cannibalization and promotional effects of counterfeits (pirated goods) on the original. In contrast, copycats and counterfeits are fundamentally different. Counterfeits copy not only the style or content but also the trademark (i.e., the brand logo), therefore they violate the trademark law. However, copycats do not copy the brand logo and are typically legal. Therefore, counterfeits can benefit the original brand by improving the awareness of the brand (i.e., promotional effect), but copycats cannot directly give consumers information about the original brands. Studies that examine the market response of copycats include Horen and Pieters (2012) and Wang et al. (2018). Horen and Pieters (2012) conducted lab experiments and survey studies at a grocery context to demonstrate how copycats can gain or lose from their resemblance to the original brands, but they remain silent about how copycats affect the demand of the original brands. Wang et al. (2018) examine the aggregate impact of copycat mobile apps on the demand of original apps. They find that deceptive and low-quality copycat apps may positively affect the demand of the original app, implying the existence of the promotional effect. In contrast, fashion goods are fundamentally different from grocery and mobile apps, in the sense that consumers mix and match multiple clothing items and peer feedback plays an important role. Moreover, as will be shown in this paper, our micro-level study specifies new mechanisms on how copycats can benefit premium brands, unlike the traditional promotional effect.

Various theoretical works have investigated fashion firms' strategies on information disclosure (e.g., Yoganarasimhan, 2012), competitive pricing (e.g., Amaldoss & Jain, 2005), given consumers' dual needs of conformity and differentiation in conspicuous consumption (Brewer, 1991). Accordingly, the consumer tradeoff between expressing individuality and conforming to others' opinions/likes, related to self and public self-consciousness for social behaviors in psychology (Fenigstein et al., 1975). Though people have heterogeneous underlying preferences, Bernheim (1994) shows that when status is very important relative to intrinsic utility, some people will conform to a single standard of behavior. To our knowledge, this is the first empirical research examining consumers' conspicuous consumption while incorporating both their intrinsic preferences and the impact of others' opinions.

To extract and quantify the styles of fashion goods, we need to analyze visual data by referring to machine learning literature for image analytics. Specifically, we apply support vector machine (SVM), support vector regression (SVR), Fast R-CNN (Girshick, 2015), and transfer learning using Siamese CNN (Hadsell et al., 2006; Veit et al., 2015) to extract clothing style features (i.e., compatibility and distinctiveness) and user appearance features (i.e., facial attractiveness and body BMI).

3.3. Data

The research context in this paper is the world's largest online fashion sharing community,[26] designed for users to post their own fashion photography, featuring themselves and their outfits. It shares similar features with other photo-based social media except that the content is restricted to fashion. More importantly, the website features a special function: a user tags the brand of the fashion items in his/her posted picture. Therefore, the brand information is clearly listed beside the fashion look[27] and can be seen by others. **Error! Reference source not found.** shows what a fashion post looks like on the website.

Figure 3.1. Example of a Fashion Post

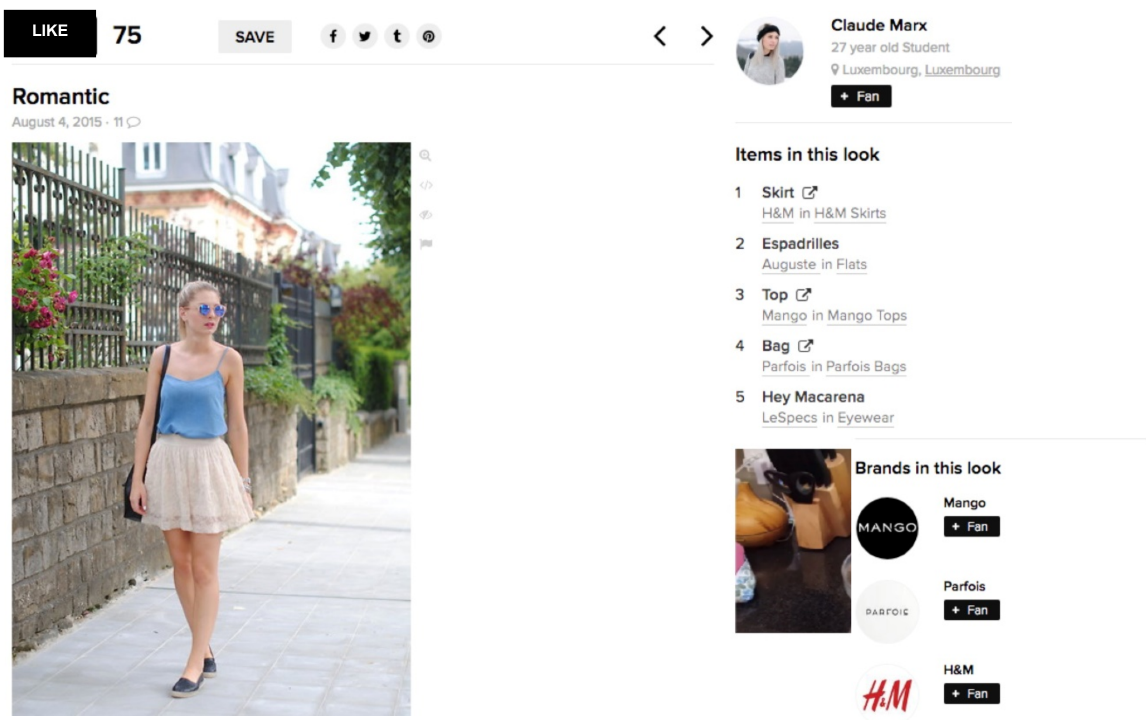

[26] The company was launched in 2008. As of July 2017, there are more than 6 million users registered.
[27] We use "fashion look" to refer to a picture.

We collect individual-level historical data from August 2013 to August 2017. The data set contains the entire history of fashion content generation for a random sample of 10,262 users[28] who registered after August 2013 and posted at least once. For each fashion look, we collect the image data, the brands for the clothing items, the time stamp, as well as how many likes the picture has attracted. For each user, we also observe his/her age from the brief biography. The gender information is not directly observable, and we will predict the gender from their picture in section 3.3.1.2.

Table 3.1. Descriptive Statistics

(a) User-level summary statistics

	#posts	#Cumulative likes	#Following	Age
mean	11.3242	173.1065	31.0892	23.4031
Std.	16.8707	1476.7562	116.6762	4.7575
min	1	0	0	4
Median	5	25	14	24
max	561	93376	6130	99

(b) Post-level summary statistics

	Mean	Std.	Min	Median	Max
#likes	25.8841	52.1091	0	10	1747

Note: "#" denotes "the count of."

Error! Reference source not found. shows the summary statistics for these users and their fashion posts. We can see that the standard deviation is large relative to the means, and the

[28] We use "fashion bloggers" and "users" interchangeably throughout the paper.

measures have skewed distribution. A small group of people has lots of posts and likes, whereas many others post very few. This observation is similar to that of most social media platforms.

Brand Categorization. Following Ha et al. (2017), we group the fashion brands into three categories: fast fashion (high street), designer, and mega couture.[29] The categorization is according to domain experts in the fashion industry, based on brand identity and price ranges. We show some examples of each brand category in Table 3.2.

Table 3.2. Brand Categorization

Brand Categories	Examples
Level 1: Fast Fashion (High Street)	Zara, H&M, Forever21
Level 2: Designer	Kate Spade, Coach, Michael Kors
Level 3: Mega Couture	Gucci, Prada, Chanel

3.3.1. Feature Extraction from Images

For a fashion look, we focus on two key aspects of visual features that can affect one's utility: the clothing styles and the appearance of the users.[30] We describe how we extract and measure the clothing styles in 3.3.1.1 and user appearance in 3.3.1.2.

3.3.1.1. Clothing Styles

In light of fashion satisfying consumers' social needs for group cohesion and differentiation (Simmel, 1904) and domain experts' opinion from well established fashion magazines, at a high

[29] The original categorization also includes "small couture" brands, but there is only one observation of such brand in our data. So we consider only three categories.

[30] On the blogging platform, the user himself/herself is the model in the picture. One account consistently posts the account owner's fashion look.

level, two style features that are particularly relevant for fashion goods: compatibility and distinctiveness. Compatibility speaks to the combination of clothing items from different categories (e.g., shirts versus pants), whereas distinctiveness measures how visually differentiated each item is from others within the same category. We abstract away from more granular style factors (e.g., color, texture) and capture the styles at a high level, because those granular factors can also be described or evaluated according to compatibility and distinctiveness. Below we explain how we extract these two style features from the fashion looks.

Figure 3.2. Steps of Clothing Style Features Extraction

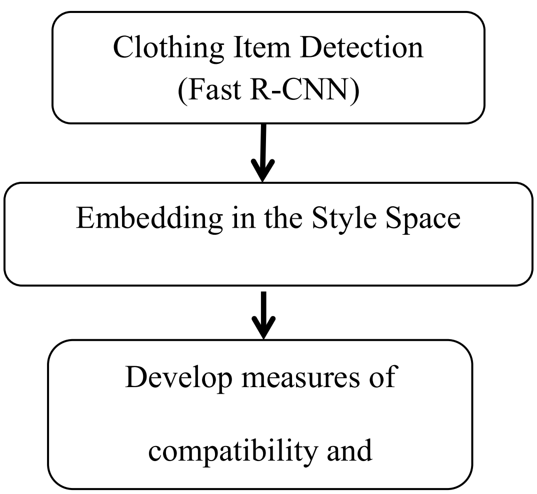

To measure fashion styles, we first need to detect or identify the clothing items in each fashion look. We follow the approach of the DeepFashion project by Liu et al. (2016). The method is based on the application of Fast Region-based Convolutional Neural Network (Fast R-CNN) (Girshick, 2015). A Fast R-CNN network takes an image and a set of object location proposals as inputs. It learns to classify objects and refine their spatial locations jointly. We adopt the network architecture of DeepFashion, which was trained on the largest and most comprehensive clothes dataset to date, annotated with clothing landmarks and categories.

For each fashion look, we extract only the clothing items, which are the most visually dominating items in a picture. Most accessories are too small to be precisely detected, so we do not include them in the analysis in this paper. We keep the cropped items (upper and bottom) if the confidence scores are higher than some threshold.[31] If the detector cannot separate the top and bottom items, we treat the clothes as full-body outfits whose compatibility is assigned an average score of the fashion looks posted during the past three months (same length as a season).

After we detect the clothing items, we can proceed to measure compatibility and distinctiveness for the fashion looks.

<u>*Compatibility*</u>

Among a large number of fashion looks, we would like to know what clothing items go well together. We adopt a deep learning approach to learn a feature transformation from images of clothing items into a latent space that represents compatibility. We use a Siamese convolutional neural network architecture (Siamese CNN) (Hadsell et al., 2006), where training data are pairs of items that are either incompatible or compatible.

To measure compatibility, we first initialize the model with weights trained on two million pairs of labeled pairs collected from the purchase data of fashion goods on Amazon.com (Veit et al., 2015). As purchasing two items together may not necessarily mean the consumers treat the items as compatible, we further collect an additional training dataset by conducting a survey on

---

[31] We tried 0.7, 0.8, and 0.9 for robustness checks.

Amazon Mechanical Turk (henceforth AMT). We directly ask survey respondents' opinion on compatibility of randomly selected pairs of items. With transfer learning, we fine-tune the deep neural network with three thousand pairs of responses (compatible versus incompatible) from the survey to improve the measure of compatibility. Please see the appendix for the survey design.

The abstraction of Siamese CNN architecture is shown in Figure 3.3. Essentially, it learns a feature transformation $f: I \rightarrow X$ from the image space I (i.e., raw representation of images) to the style space X (i.e., another representation that captures the style features). In the style space, compatible items are closer together, and incompatible items are farther away. Then, we can use the distance between two items' locations in the style space to measure how compatible they are. In Figure 3.3, I_1 and I_2 are the inputs of two clothing items from different categories (top and bottom), x_1 and x_2 are vector representations in the style space, y is the label of data (either compatible or incompatible), and θ is the set of parameters that specify the neural network, which we need to estimate.

Figure 3.3. Abstraction of Siamese CNN Architecture

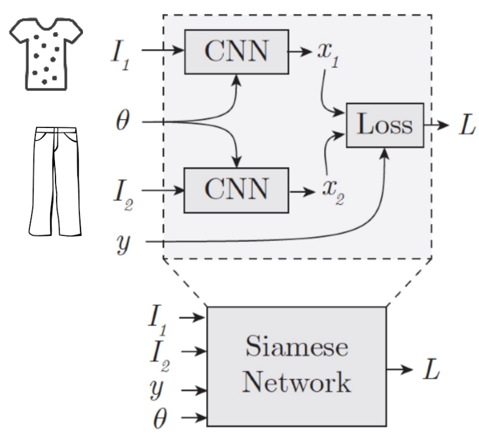

The loss function $L(\theta)$ is a contrastive loss and can be expressed as:

$$L(\theta) = \sum_{(x_1, x_2)} L_p(x_1, x_2) + \sum_{(x_1, x_3)} L_n(x_1, x_3)$$

The first term L_p penalizes when a compatible pair is too far apart, and the second term L_n penalizes when an incompatible pair is too close compared to some margin.

As illustrated in Figure 3.4, the embedding of items in the style space are vectors of dimension 256. Following Veit et al. (2015), we measure the compatibility between two items using L_2 norm. The architecture of the CNN in Figure 3.3 is based on one of the most successful network architectures, GoogLeNet (Szegedy et al., 2015), augmented with a 256-dimension fully connected layer.

Distinctiveness

We measure the distinctiveness of a clothing item by calculating how visually different the item is from all the other items in the same category. Specifically, we use the embedding in the style space to represent each clothing item's style and calculate the average style for items posted in the past three months (a season), to account for the fact that one style could be distinctive this season but may not be distinctive later on. The distinctiveness of one item is measured by L_2 norm between its style embedding and the average style. This is in a similar spirit to the creativity concept from Toubia and Netzer (2017).

Figure 3.4. Illustration of Compatibility Feature Extraction

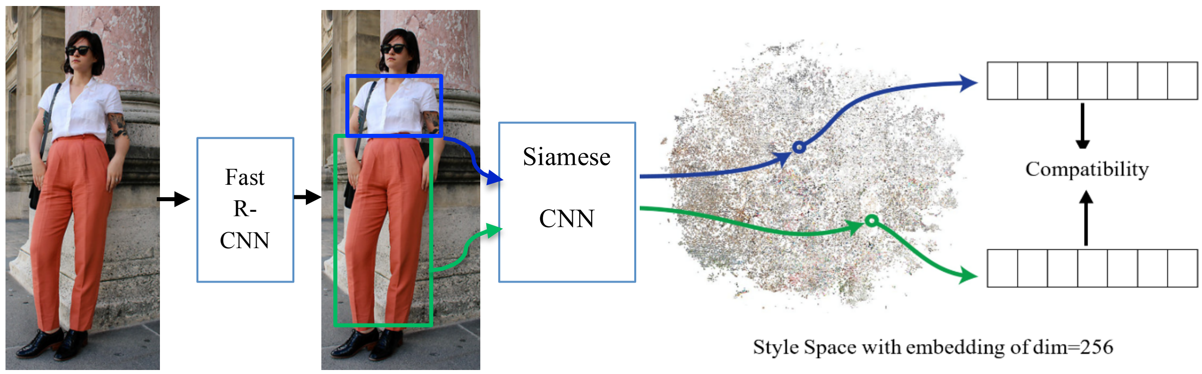

3.3.1.2. User Appearance

For the model styles, we examine the face and body features, specifically the facial attractiveness and body mass index.

Figure 3.5. Steps of Model Styles Extraction

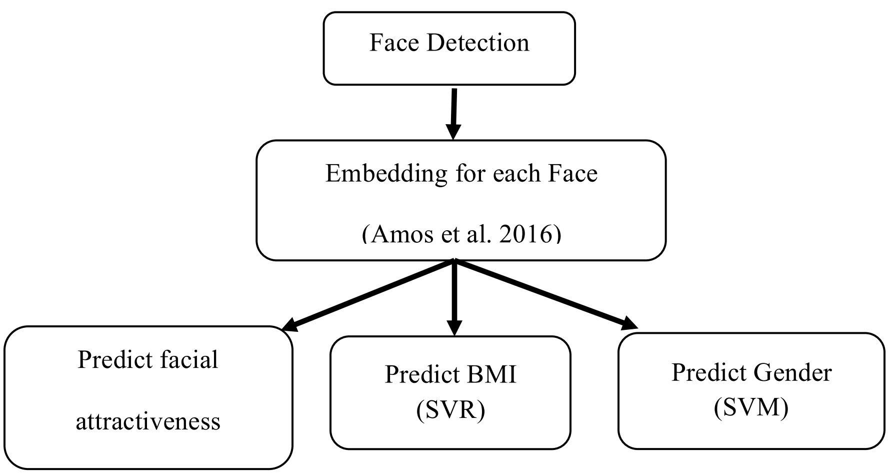

<u>Facial attractiveness</u>

The first step is to crop the face and get a vector representation of the face. We follow the deep neural network implementation by the Open Face project (Amos et al., 2016). This architecture was trained for face recognition, providing a 128-dimensional intermediate layer that represents a low dimensional embedding of any face image.

The first step is to get the low dimensional features generated using the deep neural network implementation. Then we need to train a supervised learning model to predict attractiveness. Our training data consist of three thousand images with attractiveness scored on a 1 to 7 scale, where 1 means the face is the least attractive and 7 represents the highest value of attractiveness. Each image is labeled by five Amazon mechanical turkers. We take the average of the five ratings for each image as its final rating. Given the continuous nature of the resulting rating, we train a Support Vector Regression (SVR) model that learns the relationship between the 128-dimensional image features and the attractiveness rating. The model achieves high prediction accuracy with a mean absolute error of 0.66 on the test sample.

Body feature

We measure BMI to capture the users' body feature. The training data contains 4206 images of faces with true BMI information, made available by Kocabey et al. (2017). These images are collected from Reddit posts linking to the imgur.com service. With the training data, we first crop the faces and get the embedding of 128 dimensions. Then, with the face embedding as the input and BMI (ranging from 11.5 to 50.8) as the output, we again train an SVR model that learns the relationship between the face image feature and the BMI. Eventually, we can predict BMI for a given fashion look, according to the face detected from the fashion look. The model has good performance with a mean absolute error of 2.45.

Gender information enhancement

The gender information for approximately 50% bloggers is not shown on the website. As we also want to know the gender effect, we employ SVM to predict gender from the cropped face images. As a result, we get gender information for 97% bloggers in the whole sample, with 92.89% accuracy on prediction.

3.3.1.3. Results of Feature Extraction

The objective of the Siamese CNN is to project compatible pairs close together and incompatible pairs far away. Figure 3.6 plots the distribution of distances for compatible and incompatible pairs for both before and after training for transfer learning. The plots show that the fine-tuned neural network separates the two categories with a greater margin and indicates that the network learned to separate compatible from non-compatible clothing items.

Figure 3.6. Distribution of Distances for Compatible and Incompatible Pairs

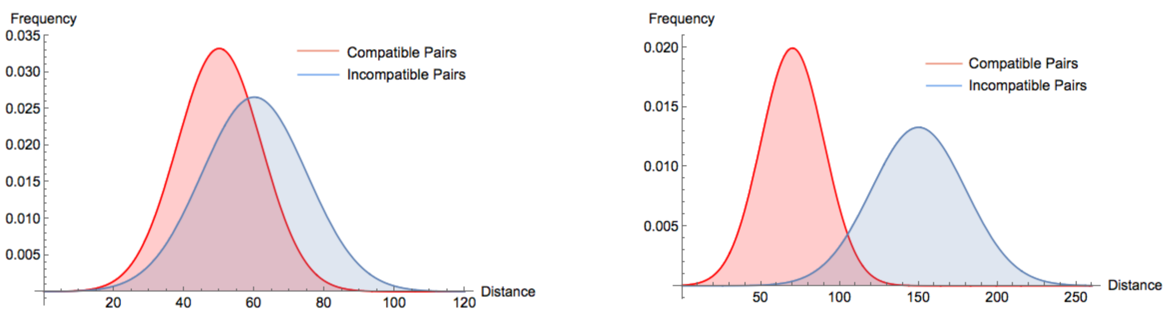

(a) Before Training (b) After Training

Table 3.3. Performance of User Features Extraction

	Facial Attractiveness	BMI	Gender
MAE	0.66	2.45	-
Accuracy	-	-	92.89%

Table 3.4. Example Photos for the Extracted Style Features

	Low Score		High Score	
Facial Attractiveness				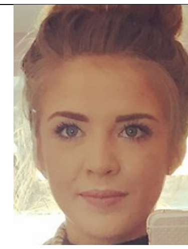
Body Feature (BMI)				
Compatibility				
Distinctiveness				

Table 3.3 reports the performance of the feature extraction tasks on facial attractiveness, BMI, and gender. For SVR task (extracting face attractiveness and BMI), the commonly used performance measure—mean absolute error (MAE)—is reported for the hold-out sample. Accuracy is reported for the binary classification tasks.

In Table 3.4, we show both high and low score examples of the extracted features, resulting from our trained learning models. Table 3.5 shows the summary statistics for the style features we extracted from the images.

Table 3.5. Summary Statistics for the Extracted Style Features

Variable	Mean	Std. Dev.	Min	Max
Facial Attractiveness	4.8201	0.6121	0.3348	7.1819
Body BMI	27.2535	3.6385	11.4966	50.7738
Compatibility	47.4943	12.4622	1.7325	257.1844
Distinctiveness	33.6818	9.8098	8.1391	124.4247

3.3.2. Exploratory Data Analysis

3.3.2.1. Users' Brand and Style Choices

After extracting the style features, we can examine the distribution of styles of fashion looks across the three brand levels to see if the style options are different for each brand level. The boxplots in Figure 3.7 and the summary statistics in Table 3.6 show that low-end brands also have pretty good styles, and the style distribution is not much different across brand levels. Therefore, for those who cannot afford luxury brands, there are always substitutes that can provide high-end styles available at lower prices.

Figure 3.7. Boxplots of Styles across Brands

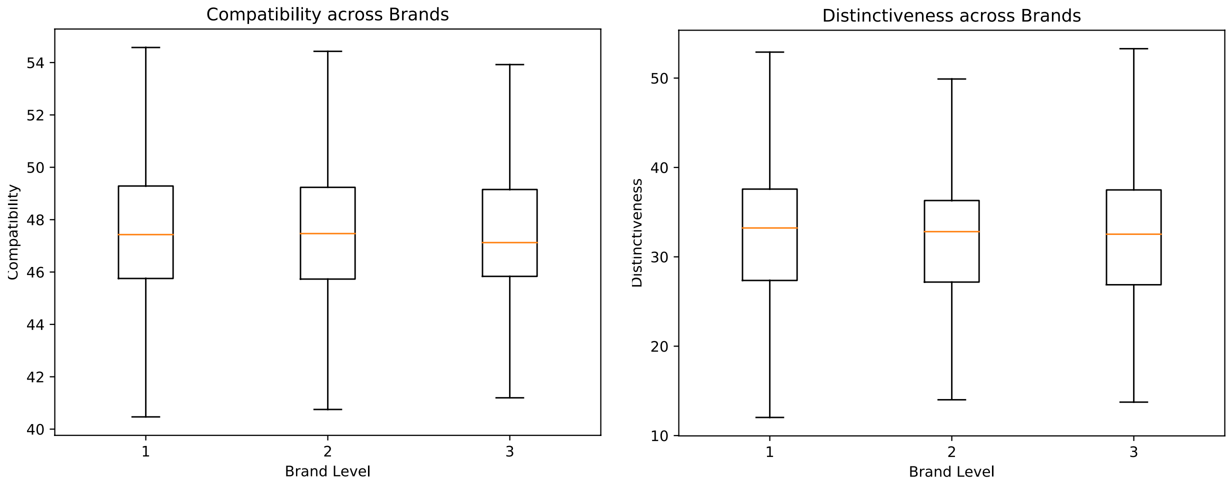

We also see from Table 3.6 that the vast majority of fashion looks contain fast fashion brands. It implies that users tend to use at least one fast fashion item to make the whole ensemble look good. Table 3.7 shows how users mix and match brands from the three brand categories. For fashion looks that adopt mega couture brands, more than 80% of them also adopt fast fashion

Table 3.6. Summary Statistics at Brand-level

	#item	face	BMI	age	gender	compatibility	distinctiveness	#likes
Fast Fashion (High Street) Brands (Adopted in 99.79% looks)								
mean	1.5652	4.8174	27.2595	24.6837	1.1875	47.4811	33.7177	30.3719
Std.	0.8540	0.6132	3.6407	5.2284	0.3903	12.4851	9.8417	59.7599
Designer Brands (Adopted in 3.05% looks)								
mean	1.0961	4.8533	27.1894	24.8999	1.2192	47.3838	32.9807	57.9064
Std.	0.2948	0.5515	3.5599	3.4180	0.4139	12.5706	9.1262	106.7643
Mega Couture Brands (Adopted in 2.00% looks)								
mean	1.1464	4.9146	26.9429	26.3511	1.1078	48.3523	33.0009	55.9045
Std.	0.3538	0.6242	3.6141	5.2377	0.3103	11.3491	9.1920	89.4968

brands. For those using designer brands, more than 86% of them mix with fast fashion brands. The data evidence shows the economic significance and importance of consumers' mixing and matching behavior in fashion choices.

Table 3.7. Mix-and-match in Fashion Looks Using Higher-end Brands

(a) For looks adopted mega couture brands

	Only Mega Couture	Mega Couture & Fast Fashion	Mega Couture & Designer	All three
Percentage	4.6%	80.15%	0.37%	14.89%

(b) For looks adopted designer brands

	Only Designer	Designer & Fast Fashion	Designer & Mega Couture	All three
Percentage	3.58%	86.40%	0.24%	9.77%

3.3.2.2. The Impact of Popularity on Choices and Inter-temporal Tradeoff

To understand how consumers choose brand and style, we need to account for a factor that may strongly affect consumers' fashion decision in our research context. That is, users are concerned about popularity or peers' likes. When a user decides whether and what to post, her level of popularity plays a big role. First, being popular can help reduce future cost, given the fact that many popular fashion bloggers are hired by fashion companies and paid to post instead of paying for what they wear. Second, the incentives for posting on social media (Lee et al., 2015) are typically social interaction and self expression, a user may derive higher utility from a given post when she is more popular and has more people watching or following. Therefore, a blogger may be strategic and forward-looking, in the sense that posting to attract others' likes to build up popularity today can help improve future utility because there would be a larger audience for future posts, as well as lower cost for future posts. In other words, when making a post decision, a blogger considers not only the current period's utility, but also the future utility. Even though posting may be worse than not posting in the current period, she may still post because it builds up popularity and the blogger can gain much more in the future.

We have explained that the discrete choice of posting is dynamic as popularity affects one's utility of posting. Moreover, a user also faces inter-temporal tradeoff when making the brand and style choices. Because the brand and style choices affect peer likes a post can attract, but the best choices for attracting likes and building up popularity may not be the choices to satisfy one's own per-period intrinsic preference. When one is not that popular, she may focus on attracting peer likes and less on self-intrinsic taste. In comparison, a popular blogger can make style and brand choices subject to less financial constraint and focus more on expressing her intrinsic fashion tastes.

An ideal measure of popularity is the number of followers. As we do not observe the number of followers for each user across time, we use the cumulative sum of likes ("SumLike," henceforth) a person has gotten from previous posts as a good proxy for his/her popularity. Figure 3.8 shows the positive linear relationship between SumLike and the number of followers at a snapshot time Aug 1st, 2017. The correlation between SumLike and the number of followers by the time Aug 2017 is 0.85, while that between the average number of likes and number of followers is 0.55. Therefore, we choose SumLike as a proxy measure for one's popularity or, more generally, the exposure one may get when posting a new fashion look.

Figure 3.8. Followers versus Likes

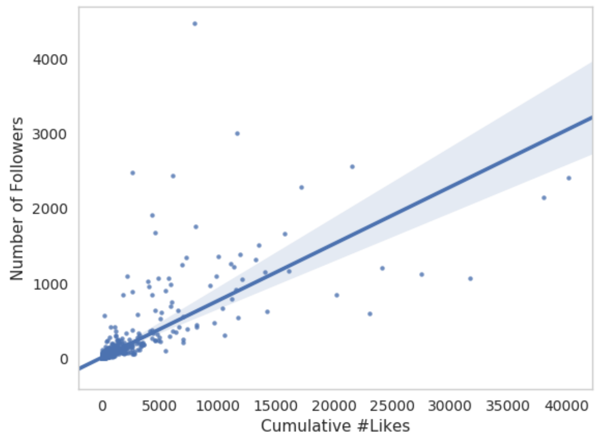

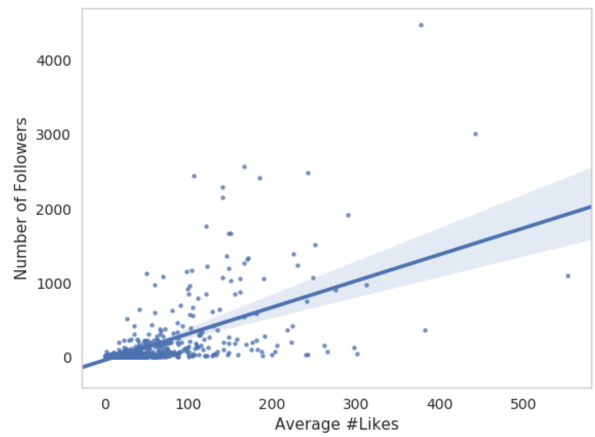

Table 3.8. Popularity Affects Decisions on Fashion Look

VARIABLES	DV: Brand 1	DV: Brand 2	DV: Brand 3	DV: Compatibility	DV: Distinctiveness
Popularity	-0.0115*	-0.00101	0.00278**	0.292***	0.0331
	(0.00609)	(0.00135)	(0.00109)	(0.0628)	(0.0449)
Observations	18,957	18,957	18,957	18,957	18,957
R-squared	0.639	0.272	0.391	0.192	0.219

Note: Fixed effects at individual level; Robust standard errors, clustered at individual level. *p<0.01;**p<0.001;***p<0.0001

In Table 3.8, we show the results of five separate regressions to see how popularity up to period $t-1$ affects one's choice of brands and styles in period t. We can see that when bloggers become more popular, they will post more mega couture brands but fewer fast fashion brands, echoing the cost decreasing effect of being popular. Moreover, they further improve the style, implying that the marginal return of improving style is higher when more people are watching.

3.3.2.3. Incentives of Posting on Social Media

People post on social media for different reasons: self-expression, social interaction, archiving, escapism, and so on (Lee et al., 2015). Below is the scatter plot of the number of posts and average likes for each user's posts. We can see that those who post more do not necessarily get more likes than users who post less. Some users' posts on average attract more than four hundred likes but the users post only around ten times, whereas other users' posts get few likes but the users post hundreds of times. Therefore, we hypothesize that the users' utility of posting is not driven only by others' likes or opinions; they may also derive utility from other channels. For example, an individual might enjoy expressing herself through posting a fashion picture, or she wants to attract those who have the same tastes. In these cases, getting likes is not the primary goal.

As a result, when modeling bloggers' decision processes, we adopt a general utility functional form, which captures not only the impact of others' likes but also bloggers' intrinsic utility derived from the brand and style choices.

Figure 3.9. Scatterplot of Posts and Likes

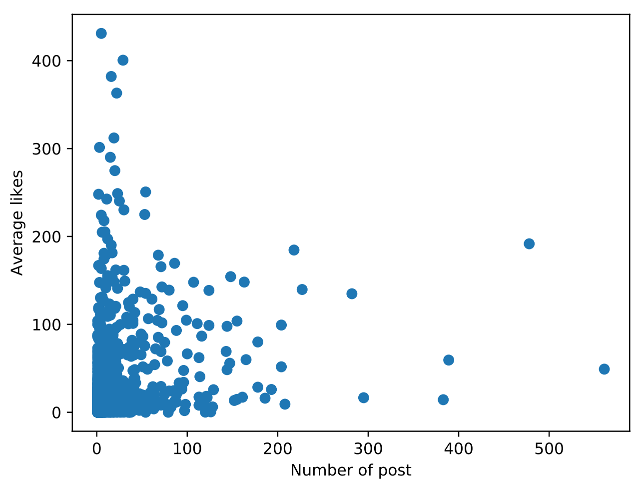

However, with the above exploratory analysis, it is still not clear how consumers value brands and styles, and how substitutable brands and styles are. Can good styles from low-end brands compensate for the utility loss from not being able to get the high-end brand with relatively worse styles? We proceed to answer this question by modeling blogger decisions at the micro-level in the next section.

3.4. Model

The timing of the events is illustrated in Figure 3.10. In each period:

1. The blogger observes the popularity she has built up with her previous posts (if there are any).
2. She decides whether to post a fashion look. If she posts, she simultaneously chooses the brand and style. If not, she goes to the next period.
3. The number of peer likes for the new post (if there are any) realizes. The next period comes.

Figure 3.10. Timing of Events

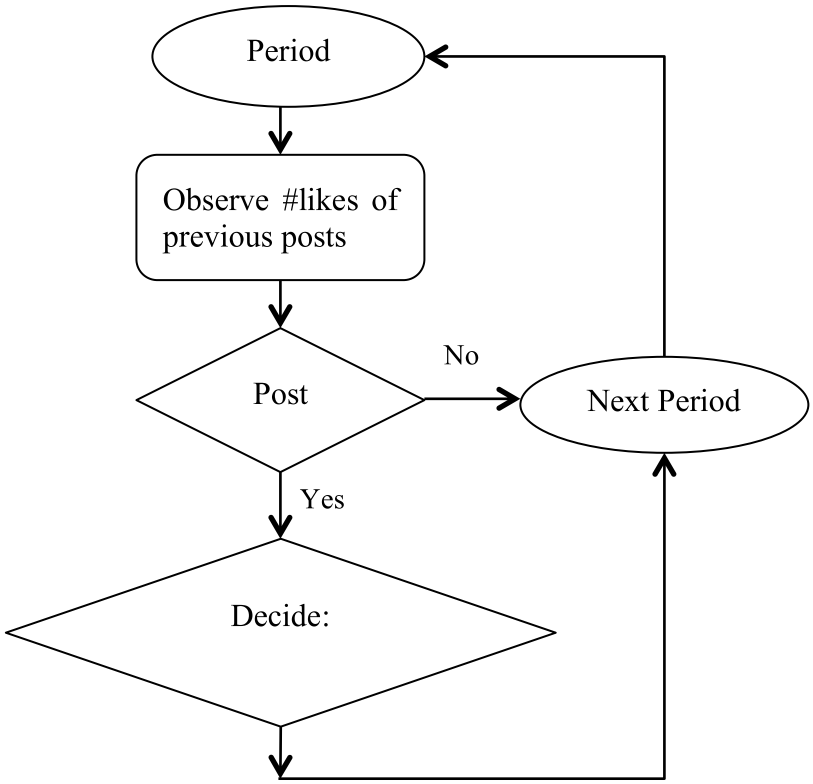

3.4.1. The Basic Model

Let St_{it}, Br_{it} denote the decision on style and brand. P_{it} is the binary decision of posting. The per-period utility of blogger i in period t follows a nested constant elasticity of substitution (CES) format (Solow, 1956):

$$u_{it}(St_{it}, Br_{it}, P_{it}|S_{it}, \theta_i) = F_{it}(S_{it}, \theta_i) \cdot R_{it}(St_{it}, Br_{it}|\theta_i) \cdot \mathbf{1}\{P_{it} = 1\} + \epsilon_{it}$$

where $R_{it}(St_{it}, Br_{it}|\theta_i)$ captures the utility gain from the fashion look attributed to the brand Br_{it} and style St_{it}.

$$R_{it}(St_{it}, Br_{it}|\theta_i) = [\alpha_{i1} {Br_{it}}^{\rho_{i1}} + (1 - \alpha_{i1}) {St_{it}}^{\rho_{i1}}]^{\frac{1}{\rho_{i1}}},$$

where the elasticity of substitution between style and brand is $r_{i1} = \frac{1}{1-\rho_{i1}}$, and α_{i1} is the share parameter. In our study, we regard the brand value as the incremental value associated with the brand name that is not related to any style attribute. For example, apart from any style features, people choose luxury brands because they work as a social label and provide hedonic rewards.

$F_{it}(\cdot)$ captures the valuation of a post. It measures the effect of popularity and individual fixed effect on the utility, apart from the brand and style choices.

$$F_{it}(S_{it}, \theta_i) = \eta_{i0} + \eta_{i1} \ln\left(1 + SumLike_{i,t-1}\right).$$

$SumLike_{i,t-1}$ is the cumulative number of likes user i got from the fashion looks that she posted in the past $t-1$ periods. In the per-period utility function, ϵ_{it} is the brand-and-style choice specific random error, following a Type-1 extreme value distribution.

3.4.1.1. Brand choice

The brand choices are discrete. The blogger makes brand choice B_{it} from categories of fast fashion (level 1), designer (level 2), and mega couture (level 3), respectively denoted by $l \in \{1,2,3\}$. For each brand level, the blogger decides how many items to include in a post. The brand choice for a fashion post is characterized by the following linear function.[32]

$$Br_{it} = exp\left(\sum_{l=1}^{3} \gamma_{i,l} x_{it,l} + \sum_{l \neq k} \gamma_{i,lk} x_{it,l} x_{it,k}\right), \qquad (2)$$

[32] We apply exp(.) to make sure the brand choice measure is positive so that it is a valid input of the CES function. If there is no brand information for a tagged item, we assume others' belief is level 1; otherwise, the user would disclose the brand information. This is consistent with the data fact that the majority of brands are level 1 demonstrated in Appendix, and conforms to consumer rationality, given higher-end brands attract more likes.

where $x_{it,l}$ denotes the number of clothing items of brand category l, in the fashion look posted by blogger i in period t. γ_{il} captures the utility gain of choosing an additional item of brand l. $\gamma_{i,lk}$ measures the utility gain of matching items from brand category l and k.

3.4.1.2. *Style choice*

The bloggers make style choices. The style, as a factor of the whole fashion look, is incorporated as a sub-nest of the CES utility function.

$$St_{it} = \left[\alpha_{i2} f_{it,1}^{\rho_{i2}} + (1-\alpha_{i2}) f_{it,2}^{\rho_{i2}}\right]^{\frac{1}{\rho_{i2}}}$$

where α_{i2} is the share parameter, and $f_{it,1}, f_{it,2}$ denote the choices of compatibility and distinctiveness. The elasticity of substitution between the two style features is $r_{i2} = \frac{1}{1-\rho_{i2}}$.

3.4.1.3. *Budget and cost*

The blogger's decisions are subject to a budget constraint y_i. Specifically,

$$\sum_{j=1}^{k} t_{ij}\, f_{it,j} + \sum_{l=1}^{3} t_{ix,l} x_{it,l} \quad \leq y_i.$$

The cost of purchasing a clothing item is allowed to change with $Like_{t-1}$.

$$t_{ix,l} = \frac{\tilde{t}_{i,l}}{1 + \delta_i \cdot \ln\left(1 + SumLike_{i,t-1}\right)},$$

where $l \in \{1,2,3\}$, and $\tilde{t}_{i,l}$ is the baseline cost for obtaining an item with brand-level l. δ_i measures the decreasing effect of the cumulative number of likes on the cost of brand choices. Bloggers may incur less cost if they are valuable to the fashion companies, as influencers can help market their products to the public. The more followers (measured by $SumLike_{i,t-1}$) a user has, the more valuable she is to the fashion company, because her posts will influence a broader audience who potentially can become customers. This is demonstrated by the fact that many fashion bloggers are paid millions every year by fashion companies.[33]

For identification purposes, we normalize the base cost for a fast fashion brand to 1—that is, $\tilde{t}_{i,1} = 1$. The budget constraint for blogger i, y_i, measures the highest cost a blogger is willing to pay for the fashion consumption, and it is assumed fixed over weeks, as the time and pecuniary resources allocated to other regular activities (e.g., working, entertainment) typically do not change much over weeks. In estimation, the budget constraint is calculated by the highest ever cost since the user started posting.

3.4.2. State Variables

The state variables are $S_{it} = \{face_i, body_i, age_i, gender_i, SumLike_{i,t-1}, \epsilon_{it}\}$, where $SumLike_{i,t-1}$ is time varying while others are fixed for the same user. With $\boldsymbol{f_{it}} = (f_{it,1}, f_{it,2})$ and $\boldsymbol{x_{it}} = (x_{it,1}, x_{it,2}, x_{it,3})$, the state transition follows

$$Like_{it} = \widehat{g}\,(SumLike_{i,t-1}, face_i, body_i, age_i, gender_i, \boldsymbol{f_{it}}, \boldsymbol{x_{it}}) + \zeta_{it}$$

where $\widehat{g}\,(\cdot)$ is estimated by a linear regression as the first step, and $\zeta_t \sim N(0, \widehat{\sigma})$. In our dataset, we observe $\hat{S}_{it} = \{face_i, body_i, age_i, gender_i, SumLike_{i,t-1}\}$. The state transition regression results are shown in section 3.6.

---

[33] http://www.blogingrace.com/highest-paid-fashion-bloggers/

3.4.3. Inter-temporal Tradeoff

Each user decides on an infinite sequence of decision rules $\{f_{it}, x_{it}, post_{it}\}_{t=1}^{\infty}$ to maximize the expected discounted utility. Substituting the brand and style choices (equation (2) and (3)) to the per-period utility function (equation (1)), we have

$$\max_{\{f_{it}, x_{it}, P_{it}\}_{t=0}^{\infty}} E_{\{S_{it}\}_{t=0}^{\infty}} \left\{ \sum_{t=0}^{\infty} \beta^t \, U_t(\boldsymbol{f_{it}}, \boldsymbol{x_{it}}, P_{it} | S_{it}, \theta_i) | x_{i0}, f_{i0}, P_{i0}, S_{i0} \right\},$$

where

$$\begin{aligned} &U_t(\boldsymbol{f_{it}}, \boldsymbol{x_{it}}, P_{it} | S_{it}, \theta_i) \\ &\quad = [\eta_{i0} + \eta_{i1} \ln(1 + SumLike_{t-1})] \\ &\quad \cdot \left\{ \alpha_i \left[exp\left(\sum_{l=1}^{3} \gamma_{i,l} x_{it,l} + \sum_{l \neq k} \gamma_{i,lk} x_{it,l} x_{it,k} \right) \right]^{\rho_{i1}} \right. \\ &\quad \left. + (1 - \alpha_i)\left[\beta_i f_{1\text{it}}^{\rho_{i2}} + (1 - \beta_i) f_{2\text{it}}^{\rho_{i2}}\right]^{\frac{\rho_{i1}}{\rho_{i2}}} \right\}^{\frac{1}{\rho_{i1}}} \cdot \mathbf{1}\{P_{it} = 1\} + \epsilon_{it} \end{aligned}$$

Let $V(S_{it})$ denote the value function:

$$\begin{aligned} V(S_{it}) = \max_{\{\boldsymbol{f_{i\tau}}, \mathbf{x}_{i\tau}, P_{i\tau}\}_{\tau=t}^{\infty}} E_{\{S_{i\tau}\}_{\tau=t}^{\infty}} &\left\{ U_t(\boldsymbol{f_{it}}, \boldsymbol{x_{it}}, P_{it} | S_{it}, \theta_i) \right. \\ &\left. + \sum_{\tau=t+1}^{\infty} \beta^{\tau - t} \, U_\tau(\boldsymbol{f_{i\tau}}, \boldsymbol{x_{i\tau}}, P_{i\tau} | S_{i\tau}, \theta_i) | x_{it}, f_{it}, P_{it}, S_{it} \right\} \end{aligned}$$

The Bellman Equation (Bellman, 1957) for the dynamic optimization problem is expressed as follows:

$$V(S_{it}) = \max_{f_{it}, x_{it}, P_{it}} E_{S_{it+1}} \{U_t(\boldsymbol{f_{it}}, \boldsymbol{x_{it}}, P_{it} | S_{it}, \theta_i) + \beta V(S_{it+1}) | \boldsymbol{x_{it}}, \boldsymbol{f_{it}}, P_{it}, S_{it}\}$$

All the decisions, that is, $\boldsymbol{f_{it}}, \boldsymbol{x_{it}}, P_{it}$, are dynamic in the sense that the current period's decisions affect the next period state through $SumLike_{t-1}$, which further affects the cost and utility of the user's brand and style choices. In other words, the users face inter-temporal tradeoffs regarding the utility derived from the post today versus its impact on the utility of posting tomorrow.

3.4.4. Heterogeneity

In the fashion sharing community, users may have different responses to others' opinions. For example, some users care a lot about others' opinion (i.e., #likes), whereas some care only about expressing themselves rather than attracting #likes. Similarly, some users base their utility heavily on the brand levels, whereas others care more about the clothing styles. Therefore, we employ a hierarchical Bayesian framework (Rossi et al., 2005) to account for heterogeneity. All structural parameters $\theta_i \in \boldsymbol{\Theta_i} = \{\alpha_{i1}, \alpha_{i2}, \boldsymbol{\rho_i}, \boldsymbol{\eta_i}, \boldsymbol{t_i}, \boldsymbol{\gamma_i}, \delta_i\}$ have random coefficient specification.[34] The prior distribution is normal for $\boldsymbol{\eta_i}, \boldsymbol{\gamma_i}$, and log normal for $\alpha_{i1}, \alpha_{i2}, \boldsymbol{r_i}, \boldsymbol{t_i}, \delta_i$, where $\boldsymbol{r_i} = \frac{1}{1-\boldsymbol{\rho_i}}$. The prior distribution is specified as $N(\mu_\theta, \sigma_\theta^2)$. We use diffuse hyper-prior distribution for all parameters.

3.5. Identification and Estimation

---

[34] The parameters in bold are vectors with subscription 1,2, etc. as specified in the model.

3.5.1. Identification

In our model, the unknown parameters include those in the state transition process ($\vec{b}$) and the primitives in the utility function ($\boldsymbol{\Theta}_i$). We briefly explain the identification of the key parameters.

First of all, the transition function parameters can be identified by the variation in the number of likes $Like_{i,t}$ and the corresponding fashion brands, style features, and blogger demographics (i.e., $SumLike_{i,t-1}, face_i, body_i, age_i, gender_i, \boldsymbol{f_{it}}, \boldsymbol{x_{it}}$). For the discount factor β , as acknowledged in the literature (e.g., Rust, 1987), it cannot be separately identified from the utility parameter, that is, the individual-fixed effect in our context. Therefore, following the conventional approach, we fixed the weekly discount at 0.998 to stay consistent with the literature (Hartmann & Nair, 2010; Liu et al., 2018).

Table 3.9. Summary of the Parameters

Notation	Explanation
α_{i1}	Share parameter for brand and style.
α_{i2}	Share parameter for different style features.
$\boldsymbol{\rho_i}$	Elasticity of substitution, $\boldsymbol{\rho_i} = \{\rho_{i1}, \rho_{i2}\}$, for brand versus style, and between style features.
$\boldsymbol{\eta_i}$	Governing the efficiency or productivity. $\boldsymbol{\eta_i} = \{\eta_{i0}, \eta_{i1}\}$, where η_{i0} captures individual fixed-effect, η_{i1} measures the effect from the cumulative likes.
$\boldsymbol{t_i}$	Cost parameters for brand and style choices. $\boldsymbol{t_i} = \{t_{i1}, t_{i2}, t_{ix2}, t_{ix3}\}$.
γ_i	Utility gain from an item of a certain brand level, $\boldsymbol{\gamma_i} = \{\gamma_{\mathrm{i}1}, \gamma_{\mathrm{i}2}, \gamma_{\mathrm{i}3}, \gamma_{\mathrm{i},12}, \gamma_{\mathrm{i},13}, \gamma_{\mathrm{i},23}\}$.
δ_i	The costing decreasing effect from cumulative #likes (as a proxy for #followers).

The elasticity of substitution between brand and style can be identified by the variation in the brand (Br_{it}) and style choices (St_{it}) across time. Similarly, the elasticity of substitution between compatibility and distinctiveness can be identified by the variation in the two style choices across periods. The brand and style choices help us back out the underlying cost of each blogger. We can separately identify the effect of $SumLike$ on the utility (η_{i1}) and on the cost (δ_i), because the cost decreasing effect is only on the brand choices, whereas η_{i1} affects both the brand and style choices in the same way. So, the different evolving patterns of brand and style choices can help us identify η_{i1} and δ_i.

In summary, the structural parameters to be estimated are $\boldsymbol{\Theta_i} = \{\alpha_{i1}, \alpha_{i2}, \boldsymbol{\rho_i}, \boldsymbol{\eta_i}, \boldsymbol{t_i}, \boldsymbol{\gamma_i}, \delta_i\}$.

3.5.2. Likelihood

The likelihood function is

$$Likelihood = L\left(\left\{\left\{\boldsymbol{f}_{it}, \boldsymbol{x}_{it}, P_{it} \middle| \hat{S}_{\mathrm{it}}, \theta_i\right\}_{t=1}^{T}\right\}_{i=1}^{I}\right) L\left(\left\{\hat{g}\left\{\hat{S}_{it} \middle| \hat{S}_{it-1}, \boldsymbol{f}_{it}, \boldsymbol{x}_{it}, P_{\mathrm{it}}\right\}_{\mathrm{t}=1}^{T}\right\}_{i=1}^{I}\right) L\left(\left\{\hat{S}_{\mathrm{i}0}\right\}_{i=1}^{I}\right)$$

where $\hat{S}_{it}$ includes all the observable states. According to the above likelihood function, the likelihood for the state transition process and that for the optimal choices can be separately estimated. Our data cover the entire history of activities for each individual, and everyone starts with $Like_{\mathrm{i}0} = 0$. As the first step, we estimate the state transition process, $\hat{g}\,(\cdot)$, with a linear regression. Then we maximize the likelihood of the optimal choices:

$$L\left(\left\{\left\{\boldsymbol{f}_{it}, \boldsymbol{x}_{it}, P_{it} \middle| \hat{S}_{it}, \theta_i\right\}_{t=1}^{T}\right\}_{i=1}^{I}\right) = \prod_{i=1}^{I} \prod_{t=1}^{T} L\left\{\boldsymbol{f}_{it}, \boldsymbol{x}_{\mathrm{it}}, P_{it} \middle| \hat{S}_{it}, \theta_i\right\},$$

where the brand choices $\boldsymbol{x_{it}}$ are discrete and style choices $\boldsymbol{f_{it}}$ are continuous.

The likelihood for each choice $\{\boldsymbol{f_{it}}, \boldsymbol{x}_{\mathbf{it}}, P_{it}\}$, consisting of both discrete and continuous choices, can be calculated through a discrete way. Note that for each combination of choice, $\{\boldsymbol{f_{it}}, \boldsymbol{x}_{\mathbf{it}}, P_{it}\}$, the predicted number of likes can be obtained with the estimated transition regression model—

that is, $\widehat{Like}_{it} = \hat{g}\,(Like_{i,t-1}, face_i, body_i, age_i, gender_i, \boldsymbol{f_{it}}, \boldsymbol{x_{it}}) \cdot \mathbf{1}\{P_{it} = 1\}$. Then, we can first eliminate all style choices that are strictly dominated. Specifically, the assumption of blogger rationality implies a unique choice set $\{\boldsymbol{f_{it}}, \boldsymbol{x_{it}}, P_{it}\}$ corresponding to the set $\{\widehat{Like}_{it}, \boldsymbol{x}_{\mathbf{it}}, P_{it}\}$. With the CES utility functional form, there exist unique closed-form solutions for $\boldsymbol{f^*_{it}}$ given $\{\widehat{Like}_{it}, \boldsymbol{x}_{\mathbf{it}}, P_{it}\}$. To put it in math, given $\{\boldsymbol{x}_{\mathbf{it}}, P_{it}\}$, the optimal choices of $\boldsymbol{f^*_{it}}$ are unique for each $\widehat{Like}_{it}$, obtained by solving the following maximization problem:

$$\max_{\{f_{ijt}\}_{j=1}^{\mathrm{k}}} \left[\beta_i f_{1\mathrm{it}}^{\rho_{i2}} + (1-\beta_i) f_{2\mathrm{it}}^{\rho_{i2}}\right]^{\frac{1}{\rho_{i2}}}$$

subject to $\hat{g}\left\{\boldsymbol{f_{it}} \middle| \boldsymbol{x}_{\mathbf{it}}, P_{it}, \hat{S}_{it}\right\} = \widehat{Like}_{it}$ and $\sum_{j=1}^{k} t_j\, f_{ijt} + \sum_{l=1}^{3} t_{xl} x_{it,l} \ \leq y_i$.

Therefore, styles choices that satisfy $\boldsymbol{f'_{it}} \neq \boldsymbol{f^*_{it}}$ but lead to the same $\widehat{Like}_{it}$ are strictly dominated, thus will never be chosen. With this observation, we can further relieve the computation burden by acting on $\widehat{Like}_{it}$ instead of multiple choices $\boldsymbol{f_{it}}$. Please see the appendix for details about how we transform the continuous choice space for $\boldsymbol{f_{it}}$ into a discrete action space on $\widehat{Like}_{it}$ while reserving the continuous nature of the style choices.

The likelihood of the optimal choices can be expressed as

$$L\left(\left\{\left\{\boldsymbol{f_{it}}, \boldsymbol{x_{it}}, P_{it} \middle| \hat{S}_{it}, \theta_i\right\}_{t=1}^{T}\right\}_{i=1}^{I}\right) = \prod_{i=1}^{I} \prod_{t=1}^{T} L\left\{\boldsymbol{f_{it}}, \boldsymbol{x}_{\mathbf{it}}, P_{it} \middle| \hat{S}_{it}, \theta_i\right\} =$$

$$\prod_{i=1}^{I} \prod_{t=1}^{T} Pr\left\{\widehat{Like}_{it}, \boldsymbol{x}_{\mathbf{it}}, P_{it} \middle| \hat{S}_{it}, \theta_i\right\} \quad .$$

With Type-1 extreme value distribution for the random error, the choice probability can be written as:

$$Pr\{\widehat{Like}_{it,\mathrm{n}}, \boldsymbol{x}_{\mathbf{it},\boldsymbol{n}}, P_{it,n}|\hat{S}_{it}, \theta_i\} = \frac{\exp\{\upsilon(\widehat{Like}_{it,\mathrm{n}}, \boldsymbol{x}_{\mathbf{it},\boldsymbol{n}}, P_{it,n}|\hat{S}_{it}, \theta_i)\}}{\sum_{n=1}^{N}\exp\{\upsilon(\widehat{Like}_{it,\mathrm{n}}, \boldsymbol{x}_{\mathbf{it},\boldsymbol{n}}, P_{it,n}|\hat{S}_{it}, \theta_i)\}}$$

where

$$\upsilon(\widehat{Like}_{it,\mathrm{n}}, \boldsymbol{x}_{\mathbf{it},\boldsymbol{n}}, P_{it,n}|\hat{S}_{it}) = U_t\left(\boldsymbol{f}_{\boldsymbol{it},\boldsymbol{n}}, \boldsymbol{x}_{\mathbf{it},\mathbf{n}}, P_{it,n}|S_{\mathrm{it}}\right) - \epsilon_{it,n},$$

and $\boldsymbol{f}_{\boldsymbol{it},\boldsymbol{n}}$ are observed from the data or backed out from $\{\widehat{Like}_{it,\mathrm{n}}, \boldsymbol{x}_{\mathbf{it},\boldsymbol{n}}\}$ if there is no post in that period. N is the total number of discrete choices. According to the data, in a fashion look, there are up to three items tagged as brand level 1, while there are up to two items tagged as brand level 2 (designer) and 3 (couture).[35] For fashion looks without brand tags, we assume audiences' belief about the brand level is fast fashion (level 1), which is consistent with the majority observation and rational behavior of tagging.[36] Therefore, there are $3\times2\times2-1=35$ brand choices if one decides to post in our context. The choice of target #likes can be from 0 to the largest #likes the bloggers have gotten on the website, which is 1747. Together with the choice of not posting, the bloggers have in total $N=35\times M+1$ discrete choices. For ease of computation, we can shrink the choice space for an individual according to the historical #likes she has gotten, plus some deviation based on $\hat{\sigma}$.

3.5.3. Estimation Methods

To estimate the infinite horizon dynamic structural model, we explore several methods, including the conditional choice probability estimation (CCP) (Hotz & Miller, 1993; Aguirregabiria &

---

[35] There are only fewer than 1% posts tagging more than the upper limits of brands; therefore, we bound the choices below 3, 2, 2 respectively for each level of brand. Though there are cases where people wear the same clothing item in multiple posts, they do not tag the same item, and there must some other new items across different posts.

[36] If the items are of higher-level brands, the consumer would tag them given the belief that no-tagged brands are level 1, because the fixed-effect regression indicates that a higher-level brand contributes positively to #likes.

Mira, 2007; Arcidiacono & Miller, 2011), the simulated method of moments (SMM) (Pakes & Pollard, 1989; McFadden, 1989), and the Bayesian estimation method (Imai, Jain, & Ching, 2009 (IJC)). SMM matches data moments with simulated moments, but it requires fully solving the dynamic optimization problem and is thus computationally very costly; moreover, it cannot capture the rich heterogeneous responses across different individuals. The CCP methods improve on the computational efficiency but can recover only very limited heterogeneity. The IJC method serves our goal to capture rich individual responses, and the computational burden is also alleviated because it requires evaluating the value function only once in each iteration.

However, IJC is designed for discrete choice models, but our model includes both discrete and continuous choices. With the modification for the dynamic choice problem explained in section 3.5.2 (please see more details in Appendix D), we can apply the IJC while reserving the continuous nature of the style choices.

3.6. Results

3.6.1. Model Comparison

We compare our model with the benchmark model without the forward looking and direct utility derived from followers. When bloggers are myopic, they do not consider the impact of today's choice on tomorrow's state, that is, number of followers. When there is no direct utility gain from followers, the only effect of followers is through decreasing cost. The utility gain from followers adds more to the model fitting than forward looking. All four measures show that the proposed model outperforms the benchmarks significantly.

Table 3.10. The Proposed Model vs. Alternative Models

	No forward-looking	No utility from followers	Proposed
Hit Rate: Post	0.564	0.532	0.877
Hit Rate: Fast Fashion	0.775	0.610	0.922
Hit Rate: Designer	0.589	0.533	0.897
Hit Rate: Mega Couture	0.520	0.421	0.811

3.6.2. Parameter Estimates

Table 3.11 shows the transition process estimated with an OLS regression. We can see that popularity does have a substantial positive effect on the peer likes for a new fashion post. The number of clothing items has increasingly positive effects with higher brand levels. More attractive faces attract more likes. The average audience likes a lower BMI, younger looks, and male models. We also see a significant positive effect of compatibility. The distinctiveness does not show a significant impact on attracting likes. The negative effect of the interaction term between high-street brands and mega couture brands shows that people do not respond favorably to outfits that match low and high-end brands together. However, the positive effect of including either a fast fashion brand or a mega couture brand dominates the negative interaction effect.

Table 3.11 reports the results for structural parameters. The estimation converges with 7000 iterations. We ran 10,000 iterations and report the results using the last 2000 iterations after burn-in.[37] Figure 3.11 shows the histogram of some structural parameters across individual

[37] Convergence was visually assessed with plots of the structural parameters. We store 100 past pseudo-value functions.

bloggers. The mean elasticity of substitution is $r_{1i} = 1/(1 - \rho_{1i}) = 1.24$ and implies that styles and brands are quite substitutable for most bloggers. In other words, they find it easy to substitute a high-end brand with good style to derive the same utility. As shown in Figure 3.11, 60.11% of the individuals treats style and brands as substitutes. For this group of consumers, high-end brands may need to worry about the copycats' cannibalization effect. For the the rest 39.89% users, they view style and brand as complements, meaning they will not be lured by only good styles, which makes the copycat problem less worrisome.

Table 3.11. Regression of State Transition

VARIABLES	OLS	Standard Error
Log (1+#Like_t-1)	12.74***	0.147
# Fast fashion (level 1)	8.814***	0.466
# Designer (level 2)	15.72**	7.237
# Mega couture (level 3)	24.18***	7.180
Face	5.768***	0.744
BMI	-0.644***	0.121
Age	-0.485***	0.0819
Gender	2.952***	0.928
Compatibility	0.151***	0.0351
Distinctiveness	0.0116	0.0448
Level 1*Level 2	-0.604	2.812
Level 1*Level 3	-6.937**	2.963
Level 2*Level 3	-1.087	6.935
Constant	-45.87***	7.532
Observations	21,093	
R-squared	0.304	

Note: "#" denotes "the count of"; Robust standard errors; Observations are collapsed at weekly level; *p<0.01, **p<0.001, ***p<0.0001.

The estimates of η_{1i} indicate that most people (about 80%) value popularity or others' attention, but there exists large heterogeneity across individuals, with a standard deviation of 0.3258. There are about 21.92% users have negative value for popularity, probably because they feel uncomfortable or stressful when more people are watching and paying attention to them. This group of users may post for archiving or escapism, which requires less or no other people's attention. From the third plot of Figure 3.11, we see there is also a large cost decreasing effect of popularity. This is consistent with the fact that influential bloggers are hired by fashion companies to promote their products.

Table 3.12. Structural Model Estimation Results

Variable	Interpretation	Mean (μ_θ)	Standard deviation (σ_θ)
α_{i1}	Share parameter	**0.5143**	0.1068
		(0.0059)	(0.0181)
α_{i2}		**0.5812**	0.0971
		(0.0054)	(0.0172)
ρ_{1i}	Elasticity of substitution	**0.1903**	0.2759
		(0.0173)	(0.0135)
ρ_{2i}		**-0.1269**	0.4131
		(0.0221)	(0.0151)
η_{0i}	Fixed effect	**0.1557**	0.4497
		(0.0265)	(0.0205)
η_{1i}	Effect of cumulative #likes	**0.4926**	0.3258
		(0.0181)	(0.0152)
t_{1i}	Cost for style choices	**1.0768**	0.4277
		(0.0263)	(0.0178)
t_{2i}		**0.9369**	0.3713
		(0.0213)	(0.0172)
$t_{x2,i}$	Cost for brand choices	**2.0601**	0.4103
		(0.0224)	(0.0143)
$t_{x3,i}$		**2.9698**	0.8729
		(0.0292)	(0.0177)
γ_{1i}	Gain from brand choices	**1.1805**	0.4763
		(0.0311)	(0.0174)
γ_{2i}		**0.6857**	0.2827
		(0.0161)	(0.0162)
γ_{3i}		**0.7848**	0.2917
		(0.0208)	(0.0178)
$\gamma_{12,i}$		-0.0130	0.3994
		(0.0212)	(0.0177)
$\gamma_{13,i}$		-0.0269	0.4013
		(0.0228)	(0.0176)
$\gamma_{23,i}$		0.0050	0.3757
		(0.0209)	(0.0146)
δ_i	Cost decreasing effect from cumulative #likes.	**0.9902**	0.3821
		(0.0220)	(0.0156)

Note: Standard errors in parentheses.

Figure 3.11. Distribution of Structural Parameters across Individuals

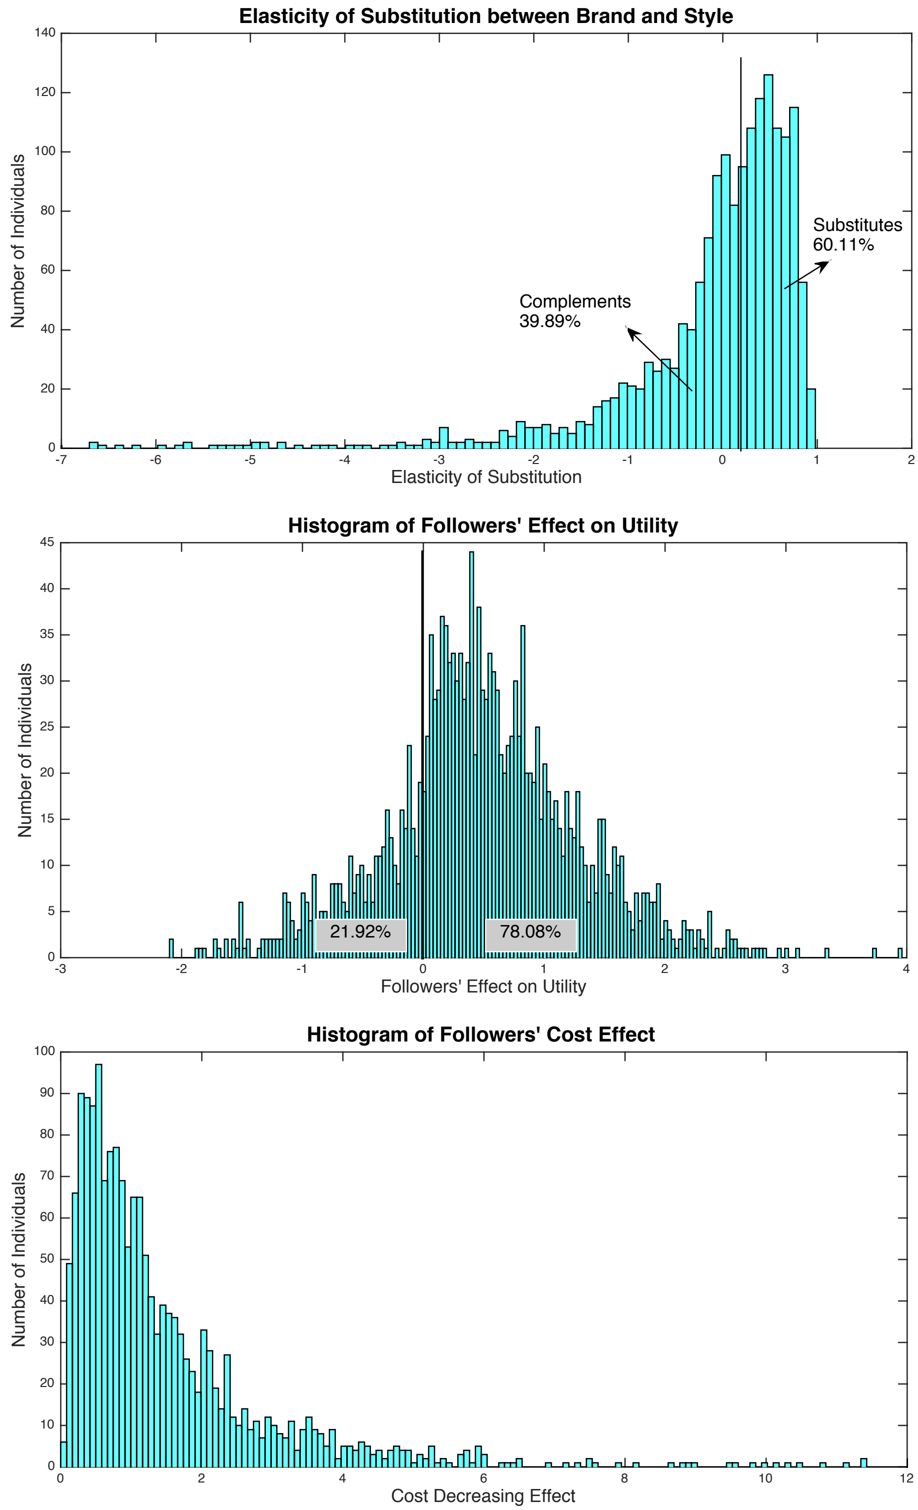

3.7. Counterfactual Studies

3.7.1. If Fast Fashion Cannot Copy Mega Couture Styles

We wanted to see what would happen to consumers' choices if copyright law provided more protection for fashion designs. In this counterfactual world, the fast fashion companies are prohibited from producing styles similar to the original fashion styles of mega couture brands. As a result, the style choices at the fast fashion level would be more restricted. Specifically, the styles of the high-end brands become relatively more exclusive or distinctive, whereas the styles at the fast fashion level become not distinctive, according to our measurement of distinctiveness. Moreover, as fast fashion companies cannot produce copycat styles in the counterfactual world, if consumers wear both the high-end and low-end styles together, the compatibility would be much lower than before. How would exactly the style options change? For each mega couture item, we first calculate its distance to all other fast fashion items, using the vector representation. We experimented by dropping the fast fashion styles with 10% and 5% smallest distance (largest similarity) and replace the removed ones with the average of the other fast fashion styles. Then we calculate the distinctiveness for the fast fashion styles, and the highest compatibility between a fast fashion and a higher-end brand. We observe the remaining fast fashion styles are below about 25 percentile of the mega couture styles. Therefore, we operationalize the analysis by restricting the distinctiveness of the fast-fashion styles to be bounded below 25 percentiles of the styles from the mega couture brands. We also restrict the compatibility of matching fast fashion with mega couture brands to be within the bottom quartile of other combinations.

The first column of Table 3.13 shows the results of a hundred simulations. On average, the posting probability drops by 7.92%. Interestingly, not only the fast fashion brands are worse off, but also the adoption of all brands decreases. Why are the high-end brands also worse off when copycats are prohibited? By comparing consumers' choices in the counterfactual world and what they chose before, we found three mechanisms contributing to a lower demand for the mega couture brands.

First, many bloggers combine clothing items across brand levels to make a complete outfit. When the style choice is restricted at low-end brands, consumers are subject to higher financial pressure to buy high-end brands to get a satisfactory ensemble. In other words, this is driven by the consumers who cannot always afford high-end brands for their ensembles and who therefore mix and match both low-end and high-end brands. The counterfactual policy would put this group of consumers unsatisfied with the styles of mixing and matching high-end and low-end brands together, resulting in them buying nothing. Thus, they end up buying fewer clothes and post less. The data facts (Table 3.7) suggest that about 80% (86%) of fashion looks incorporating a mega couture (designer) brand also include items at the fast-fashion level, further demonstrating the economic significance of the mix-and-match mechanism in the fashion market.

Second, some consumers may not value high-end brands that much. But they can accumulate popularity across time by wearing attractive styles from fast fashion. Once their popularity is high, there are more people following and paying attention to them, then they will derive more value from what they wear, including high-end brands. Therefore, they will be more likely to adopt high-end brands.

Third, many consumers cannot afford mega couture brands in the very beginning, so they rely on trendy styles from the low-end brands to build up popularity. Once the popularity is high enough, they start to be able to afford more high-end brands due to the cost decreasing effect of popularity. The three above-mentioned mechanisms are market expansion effect brought by fashion copycats. Our findings demonstrate that copycats' market expansion effect dominates the competition effect, leading to a positive net effect on the demand of high-end brands. In addition, the results also show a boost in the choice of compatibility, which for many bloggers is a substitute for distinctiveness.

Discussion on Firm Strategy

As we do not have firm side data and only model the consumers' decision making, our results speak to the cases where fast fashion firms do not find a way to dramatically change their styles and firms charge the original price. This section discusses how firms would react to the counterfactual policy and the corresponding demand effect on high-end brand.

Fast fashion firms would either put the same effort in designing their own styles, which is less distinctive and matches not well with higher-end brands, but decrease the price to attract more consumers. This escalated price competition would further decrease the demand for mega couture brands. Alternatively, the low-end brands may invest more in coming up a large number of their own styles to make them not only distinctive but also compatible with higher-end brands. We keep ignorant about how they can achieve this goal with the current low price, but one possibility is that they may transform into designer brands, providing better styles but also higher price due to the higher production cost. The demand effect of prohibiting copycats on mega couture brands is most likely still negative because there is only a small proportion of consumers mix and match designer brands with mega couture brands.

Regarding mega couture brands' strategic reaction, they may start a low-end product line, providing similar styles to those from high-end product line but charge lower price (i.e., umbrella branding). However, in this case, the low-end product line may erode the parent brand's value, and the demand effect is not clear without additional information and further empirical study. The other strategy the mega couture brands may use is to simply lower the price thus attract more demand. In this case, the mega couture brands' profit is definitely lower than before because they have to charge a lower price to achieve the same demand.

Overall, we can see if fast fashion copycats were prohibited, neither fast fashion firms nor mega couture brands can find an easy way to combat the loss.

Discussion on Generalizability

As we acknowledged upfront, an average consumer may differ from a fashion social media user. To generalize to all other consumers who do not use social media or do not post fashion content on social media, we need further study and additional data which is not available for now. However, we still can try to make the results more generalizable. In our sample, there are around 10% users are professional fashion bloggers who operate their own blogging website. They may have very different incentive than an average consumer. So we exclude these professional bloggers and check the choice change for other users. The results are reported in the second column of Table 3.13. Compared to the first column, we see the net effect is still negative but with smaller magnitude. This is due to the fact that professional bloggers care more about popularity, and the lack of fast fashion copycats makes it harder to build up popularity. As a result, the value increasing effect and cost decreasing effect of popularity cannot work their way to increase adoption of mega couture brands.

3.7.2. Offline Market: Absence of Follower Effect

This provides a similar scenario as the offline fashion market, where people do not have the "likes" information for what they wear. The posting decision is analogous to the purchase decision. There are two consequences. First, one's utility does not involve the likes' impact; second, the cost of purchasing will not decrease over time.

We seek to evaluate how less/more likely a blogger would post. This analysis allows us to understand how many more purchases can be achieved by the existence of social media, compared to the traditional offline market. Companies can leverage the "follower function" of social media to boost their revenue. The results in the fourth column of Table 3.13 show that there would be a 9.07% drop in post probability if bloggers did not know about others' following. High-end brands suffer most with the largest drop in choice probability, because the cost decreasing effect is also deprived, and this affects the purchase of high-end brands most.

3.7.3. When Brands Are Not Observable

As previously mentioned, our data context has a special brand tagging function, which requires users to pinpoint the brands for the clothing items in their fashion look and clearly lists the brands besides the picture. In most other social medium, there is no such feature. Moreover, in the offline market, brand logos are usually hidden or not obvious. In this counterfactual study, we investigate how fashion bloggers make choices if there were no way to inform others about the brand information. Consistent with the main analysis, when there is no brand information for a clothing item, the belief is that it is a high street brand (i.e., low-end brand). The results are shown in the third column of Table 3.13.

Table 3.13. Results of Counterfactual Studies

	Counterfactual Policies				
	No copycats	No copycats (Sub sample)	No brand info	No followers	Targeted ranking
Δ#Posts (%)	-7.92%	-6.12%	-0.21%	-9.07%	0.90%
Δ#Fast fashion (%)	-5.78%	-5.52%	0.74%	-6.07%	0.61%
Δ#Designer (%)	-7.68%	-6.08%	-0.25%	-11.78%	0.87%
Δ#Mega couture (%)	-12.44%	-9.32%	-0.30%	-19.86%	0.89%
ΔCompatibility (%)	-3.34%	-3.51%	99.93%	-	-
ΔDistinctiveness (%)	-23.87%	-15.76%	19.88%	-	-

The results show that there would be 0.21% fewer posts. The users would post 0.74% more fast fashion brands, but 0.25% fewer designer brands and 0.30% fewer mega couture brands. Moreover, both style features are improved. Intuitively, when higher-end brands cannot be identified, bloggers cannot use high-end brands to attract peer likes anymore. Therefore, for those who care about popularity and peer likes a lot, they have fewer means to achieve the goal, resulting in lower utility from blogging and thus fewer posts. On the other hand, they switch focus to either using more low-end brands or investing more on styles. As the cost of a low-end brand is lower, bloggers would be able to purchase more low-end clothes, resulting in more posts of fast fashion brands. In summary, for higher-end brands, hiding or obscuring the brand logo would benefit the lower-end brands but hurt high-end brands, given all brands can provide similar styles. This result implies that, high-end brands may want to design their brand logos obvious enough to consumers, so that those who care about peer likes would have greater incentive to buy.

3.7.4. Platform Ranking System: If "Sensitive" Users Are Prioritized

As bloggers value others' opinions differently, the website can change the ranking system so that it favors those "sensitive" bloggers, to incentivize more fashion posts. Effectively, the platform can exogenously change the number of likes a post can get through the ranking system. We run a counterfactual analysis on the top 10% of sensitive bloggers by exogenously giving one more like to their posts. The results in the last column of Table 3.13 show that, on average, there would be a 0.90% increase in the probability of posting in each period. There is a 0.61%, 0.87%, and 0.89% increase in the number of items posted from brand levels 1, 2, and 3 respectively.

3.8. Contribution and Limitations

In the fashion market, it has been argued that the fast fashion brands copy the designs of the high-end fashion brands. This practice can potentially reduce the distinctiveness of the luxury fashion brands thus erode their brand equity. However, there is no systematic study attempting to investigate the impact of fast fashion copycats on high-end brand equity and the underlying reasons. The key challenges limiting such study are the lack of scalable ways to quantify fashion styles and the unavailability of large-scale data on individuals' choices over brands and styles. In this paper, we use the user-generated data a large fashion sharing platform and the state-of-the-art deep learning methods on image analytics to quantify fashion styles. Given fashion social media users' significant impact on fashion trend and demand, understanding their decision process sheds light on fashion consumers' choices across population. For fashion goods, brand

and style are two of the product attributes consumers care about most. Incentive-wise, fashion bloggers make the tradeoff between self-intrinsic tastes and building up popularity. To figure out the underlying reasons of how fast fashion may affect high-end brands' demand, we build a structural model to investigate fashion social media users' decision processes that reveals their heterogeneous responses to brands, styles, and popularity.

Our results show that styles and brands are quite substitutable for most people. In other words, they find it easy to substitute a high-end brand with good styles to derive the same utility. These are the consumers that high-end brands could lose to the low-end brands providing comparable styles. We also find that most users value being popular (or peer likes), but there exists a large variance in how much they value popularity. This variation explains why some people keep posting even though they get almost zero peer likes all the time. Moreover, we find that a higher popularity can help reduce posting cost which is consistent with the fact that fashion bloggers with lots of followers are sponsored by fashion companies to post about their products.

In the main counterfactual analysis, we restrict the availability of style choice for fast fashion brands. The results show that not only the fast fashion brands will suffer, but also the high-end brands will be worse off. This means that a more restrictive copyright law on fashion design may not necessarily help the mega couture brands. We found three mechanisms that contribute to this result. Because the more restricted choice of styles from low-end brands would limit the mix-and-match choices for consumers and put them on greater financial constraint to get a satisfactory ensemble of clothes, resulting in them buying less high-end brands. Moreover, the lack of good styles from low-end brands will make it harder for consumers to build up popularity/likeability, which limits the value of what they wear, including high-end brands. A low popularity also makes it unlikely to get a cost reduction from high-end brands. All these reasons indicate that copycats can benefit the high-end brand, as our findings suggest the market expansion effect dominates the competition effect. We also simulate the case where brand information cannot be seen, which is similar to the offline market where brand logos are hidden or other social media where brand tagging is not featured. The results indicate that on average, fast fashion benefits, whereas the high-end brands suffer. Another counterfactual analysis indicates that there would be a 9.07% drop in post probability if users did not know about peer likes, and high-end brands would suffer the most. This scenario is similar to that of the offline market. This analysis, therefore, allows us to understand how many more purchases can be achieved by the mere existence of social media, compared to the traditional offline market. We also find that an alternative ranking system that prioritizes people who more highly value "likes" can hugely increase the amount of content generated on the platform.

There are several limitations of the current paper that call for future studies. First, the style measure is an objective measure, whereas different consumers may evaluate the compatibility and distinctiveness differently. So the interpretation of the results should be based on the objective style measure we use. An ideal scenario is that we get to know how each consumer evaluates styles. Though this is not feasible, future research could do a more targeted analysis if there were more granular information about consumers from lab experiments or surveys and could then match people's evaluations of styles based on these observables. Second, though we tried to capture the most salient factors in a fashion post—that is, the model face, the body feature, and the style of clothing items—we did not include everything that may affect the likes a post can get. For example, the accessories may matter. The challenge is that although the detection algorithms for fashion items are state-of-the-art, they are not perfect, especially for items that are too small (e.g., earrings and hats are comparatively quite small). Had we been able to use a better detection algorithm or rich training set, we could have incorporated more factors in the model.

Reference

Aguirregabiria, B Y Victor, and Pedro Mira. 2007. “Sequential Estimation of Dynamic Discrete Games.” *Econometrica* 75 (1): 1–53.

Amaldoss, Wilfred, and Sanjay Jain. 2005. “Conspicuous Consumption and Sophisticated Thinking.” *Management Science* 51 (10): 1449–66.

Amaldoss, Wilfred, and Sanjay Jain. 2005. “Pricing of Conspicuous Goods: A Competitive Analysis of Social Effects.” *Journal of Marketing Research* 42 (1): 30–42.

Amos, Brandon, Bartosz Ludwiczuk, and Mahadev Satyanarayanan. 2016. “OpenFace: A General-Purpose Face Recognition Library with Mobile Applications.” *Technical Report CMU-CS, CMU School of Computer Science* 16 (118).

Arcidiacono, Peter, and Robert Miller. 2011. “Conditional Choice Probability Estimation of Dynamic Discrete Choice Models With Unobserved Heterogeneity.” *Econometrica* 79 (6): 1823–67.

Bernhein, Douglas. 1994. “A Theory of Conformity.” *Journal of Political Economy* 102 (5): 841–77.

Borkovsky, Ron N., Avi Goldfarb, Avery M. Haviv, and Sridhar Moorthy. 2017. “Measuring and Understanding Brand Value in a Dynamic Model of Brand Management.” *Marketing Science* 36 (4): 471–99.

Brewer, Marilynn. 1991. "The Social Self: On Being the Same and Different at the Same Time." *Personality and Social Psychology Bulletin* 17 (5): 475–82.

Fenigstein, A, M. F Scheier, and A.H Buss. 1975. "Public and Private Self Consciousness: Assessment and Theory." *Journal of Consulting and Clinical Psychology* 43 (4): 522–27.

Girshick, Ross. 2015. "Fast R-CNN." *arXiv:1504.08083v2 [cs.CV]*.

Goldfarb, Avi, Qiang Lu, and Sridhar Moorthy. 2009. "Measuring Brand Value in an Equilibrium Framework." *Marketing Science* 28 (1): 69–86.

Ha, Yu-I, Sejeong Kwon, Meeyoung Cha, and Jungseock Joo. 2017. "Fashion Conversation Data on Instagram." *International AAAI Conference on Web and Social Media. arXiv:1704.04137v1 [stat.ML]* 13 Apr.

Hadsell, Raia, Sumit Chopra, and Yann LeCun. 2006. "Dimensionality Reduction by Learning an Invariant Mapping." *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* 2: 1735–42.

Hartmann, Wesley R., and Harikesh S. Nair. 2010. "Retail Competition and the Dynamics of Demand for Tied Goods." *Marketing Science* 29 (2): 366–86.

Helmore, Edward. 2010. "Google Targets Fashion Market." *Guardian*.

Horen, Femke van, and Rik Pieters. 2012. "When High-Similarity Copycats Lose and Moderate-Similarity Copycats Gain: The Impact of Comparative Evaluation." *Journal of Marketing Research* XLIX (February): 83–91.

Hotz, V Joseph, and Robert a Miller. 1993. "Conditional Choice Probabilities and the Estimation of Dynamic Models." *The Review of Economic Studies* 60 (3): 497–529.

Hu, Y., L. Manikonda, and S. Kambhampati. 2014. "What We Instagram : A First Analysis of Instagram Photo Content and User Types." *Proceedings of the Eight International AAAI Conference on Weblogs and Social Media*.

Imai, Susumu, Neelam Jain, and Andrew Ching. 2009. "Bayesian Estimation of Dynamic Discrete Choice Models." *Econometrica* 77 (6): 1865–99.

Kamakura, Wagner A, and Gary J Russell. 1993. "Measuring Brand Equity with Scanner Data." *International Journal of Research in Marketing* 10 (1): 9–22.

Keller, Kevin Lane, and Donald R. Lehmann. 2006. "Brands and Branding: Research Findings and Future Priorities." *Marketing Science* 25 (6): 740–59.

Kocabey, Enes, Mustafa Camurcu, Ferda Ofli, Yusuf Aytar, Javier Marin, Antonio Torralba, and Ingmar Weber. 2017. "Face-to-BMI: Using Computer Vision to Infer Body Mass Index on Social Media." *arXiv:1703.03156v1 [cs.HC]*.

Kuksov, Dmitri. 2007. "Brand Value in Social Interaction." *Management Science* 53 (10): 1634–44.

Liu, Liu, Daria Dzyahura, and Natalie Mizik. 2018. "Visual Listening In: Extracting Brand Image Portrayed on Social Media." *Working Paper*.

Liu, Ziwei, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. 2016. “DeepFashion: Powering Robust Clothes Recognition and Retrieval with Rich Annotations.” *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, no. 1: 1096–1104.

Ma, Liye, Alan L Montgomery, and Michael D Smith. 2016. “The Dual Impact of Movie Piracy on Box - Office Revenue : Cannibalization and Promotion.” *Working Paper*.

Nam, Hyoryung, Yogesh V. Joshi, and P.K. Kannan. 2017. “Harvesting Brand Information from Social Tags.” *Journal of Marketing* 81 (4): 88–108.

Oberholzer-Gee, Felix, and Koleman Strumpf. 2007. “The Effect of File Sharing on Record Sales: An Empirical Analysis.” *Journal of Political Economy* 115 (1): 1–42.

Qian, Yi. 2008. “Impacts of Entry by Counterfeiters.” *The Quarterly Journal of Economics* 123 (4): 1577–1609.

Qian, Yi. 2014. “Counterfeiters: Foes or Friends? How Counterfeits Affect Sales by Product Quality Tier.” *Management Science* 60 (10): 2381–2400.

Simmel, Georg. 1904. “Fashion.” *International Quaterly* 10 (22): 130–55.

Smith, M D, and R Telang. 2009. “Competing with Free: The Impact of Movie Broadcasts on DVD Sales and Internet Piracy.” *MIS Quarterly* 33 (2): 321–38.

Szegedy, Christian, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, et al. 2015. “Going Deeper with Convolutions.” *CVPR*.

Toubia, Olivier, and Oded Netzer. 2017. “Idea Generation, Creativity, and Prototypicality.” *Marketing Science* 36 (1): 1–20.

Veit, Andreas, Balazs Kovacs, Sean Bell, Julian Mcauley, Kavita Bala, and Serge Belongie. 2015. "Learning Visual Clothing Style with Heterogeneous Dyadic Co-Occurrences." *arXiv:1509.07473v1 [cs.CV]*.

Wang, Quan, Beibei Li, and Param Vir Singh. 2018. "Copycats versus Original Mobile Apps: A Machine Learning Copycat Detection Method and Empirical Analysis." *Information Systems Research* 29 (2): 273–91.

Yoganarasimhan, Hema. 2012. "Cloak or Flaunt? The Fashion Dilemma." *Marketing Science* 31 (1): 74–95.

Appendix

A. More on summary statistics

Table A.14 shows how many times each brand appears in the posts in total.

Table A.14. Frequency of brands (Top 60)

Brand	Count	Brand	Count	Brand	Count	Brand	Count
H&M	35891	Dr. Martens	2743	GUESS	1490	Celine	895
Zara	28865	Topman	2546	Adidas	1393	Prada	876
Forever 21	15384	Stradivarius	2376	Diy	1382	Choies	847
Topshop	10273	River Island	2325	Cheap Monday	1359	Blanco	793
Vintage	8493	Aldo	2265	Shein	1306	Gucci	788
Mango	5518	Pull&Bear	2169	Uniqlo	1288	J. Crew	783
Primark	5504	Nike	2160	Marc by Marc Jacobs	1258	Ralph Lauren	776
Asos	5329	Thrifted	2051	Louis Vuitton	1156	C&A	772
American Apparel	5006	Jeffrey Campbell	2040	Cotton On	1139	Vero Moda	768
Converse	3922	Vans	1978	New Yorker	1106	Old Navy	755
Ray-Ban	3716	Chanel	1836	Gina Tricot	1053	Nine West	749
Levi's(r)	3597	Gap	1644	Monki	957	Calvin Klein	742
Urban Outfitters	3450	Steve Madden	1575	OASAP	908	Pimkie	732
New Look	3271	Romwe	1564	Target	902	Alexander Wang	732
Bershka	3169	Michael Kors	1523	FrontRowShop	901	Saint Laurent	671

B. Survey

We are interested in assessing the aesthetic quality of human photos on three factors: clothing compatibility, facial attractiveness and body attractiveness. These are subjective measures, but you can guide your judgments according to the following instruction and examples.

Instruction.

1. Clothing Compatibility: how well the clothing items match together.

Example.

The 1st item is compatible with the 2nd item but not compatible with the 3rd item.

2. Facial (Body) attractiveness is the degree to which a person's facial (body) features are considered aesthetically pleasing or beautiful. Below are some general aspects you can refer to.

<table>
<tr><th>Facial attractiveness</th><th>Body attractiveness</th></tr>
<tr><td>
· Symmetry. Symmetric images are usually viewed as more attractive than asymmetric images. For example, (a) is more attractive than (b) in the following picture.


· Beauty of the average. A face is usually considered more attractive if it more closely resembles the majority of other faces within a population; non-average faces have more extreme characteristics than the average of a population. Examples of average faces are shown below:
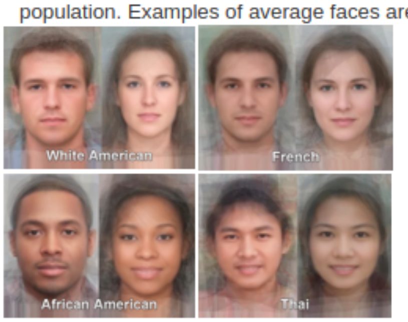

· Facial health (skin health and color). High healthiness is typically associated with higher score of attractiveness. As shown in the example below, the right face is more attractive than the left face.

</td><td>
· Symmetry. Similar to a face, a symmetric body is viewed as more attractive than an asymmetric one.
· Body Health. Healthy bodies are considered more attractive. For example, overweight or emaciated bodies are considered of low attractiveness.

All the following aspects also affect a body's attractiveness. However, preference varies across different cultures and regions.

· Leg-to-body ratio (LBR). In the following pictures for both female and male, the LRB of the left body is lower.
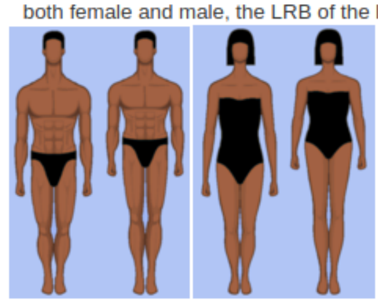
· Waist-to-hip ratio.
· Breast-to-waist ratio.
· Height.
</td></tr>
</table>

Examples.

Below we offer some examples of faces and bodies with different attractiveness levels (1 to 10). Please look at these examples to set your expectations.

Examples of faces with various attractiveness levels in randomized order (1st row for women, 2nd row for men)

Examples of bodies with various attractiveness levels in randomized order (1st row for women, 2nd row for men)

Questions.

[Picture given here]

Please answer the following questions regarding the given picture. For the first two questions, please disregard clothing styles and focus only on face (body) attractiveness:

[Compatibility] Do you think the clothing items in the photo are compatible (match well)?
Yes/No

[Body attractiveness] How attractive do you think the person's **body** is? (1: least attractive, 10: most attractive)
(Scale 1~10)

[Facial attractiveness] How attractive do you think the person's **face** is? (1: least attractive, 10: most attractive)

(Scale 1~10)

[Picture aesthetics] How visually pleasing do you think the whole picture is? (1: least visually pleasing, 10: most visually pleasing)
(Scale 1~10)

[Clothing Fashion] How fashionable is the clothing style? (1: worst fashion, 10: best fashion)
(Scale 1 ~10)

[Clothing Price Appearance] How expensive do you think the clothings are? (1: very cheap, 10: very expensive)
(Scale 1 ~10)

[Subject Gender] Is the person in the picture a female or male?
Male/Female

[Subject Glasses] Is the person wearing glasses in this picture?
Yes/No

What is **your** age?
A: <20
B: 20~29
C: 30~39
D: 40~49
E: 50~60
F: >60

What is **your** gender?
Female/Male

What is **your** ethnicity?
A) White
B) Black or African American
C) East Asian (e.g., Chinese, Korean, Japanese)
D) Other Asian
E) Hispanic or Latino
F) Other

C. Solve for optimal style choices

For a given targeted $\widehat{Like}_t$, which affects the future utility, each individual would choose styles that maximize their current utility. That is, user i solves the following optimization problem (ignoring the subscript i):

$$\max_{(f_1, f_2)} [\beta f_1{}^{\rho_2} + (1 - \beta) f_2{}^{\rho_2}]^{\frac{1}{\rho_2}}$$

s.t. $b_1 f_1 + b_2 f_2 = \mathrm{L}$

where L is a constant given by

$$L = \widehat{Like}_t - \vec{\mathbf{b}} \cdot (x_1, x_2, x_3, x_1 x_2, x_1 x_3, x_2 x_3, Like_{t-1}, face, body, age, gender),$$

and $\vec{\mathbf{b}}$ is the coefficients estimated from the state transition regression $\widehat{g}\,(\cdot)$.

For the Lagrangian

$$H = [\beta f_1{}^{\rho_2} + (1 - \beta) f_2{}^{\rho_2}]^{\frac{1}{\rho_2}} + \lambda[L - (b_1 f_1 + b_2 f_2)]$$

The first order conditions are

$$L_{f_1} = \frac{1}{\rho_2}[\beta f_1{}^{\rho_2} + (1 - \beta) f_2{}^{\rho_2}]^{\frac{1}{\rho_2}} \beta \rho_2 f_1{}^{\rho_2 - 1} - \lambda b_1 = 0$$

$$L_{f_2} = \frac{1}{\rho_2}[\beta f_1{}^{\rho_2} + (1 - \beta) f_2{}^{\rho_2}]^{\frac{1}{\rho_2}} (1 - \beta) \rho_2 f_2{}^{\rho_2 - 1} - \lambda b_2 = 0$$

Solving the above system of equations, we have

$$\frac{b_1}{b_2} = \frac{\beta f_1{}^{\rho_2 - 1}}{(1 - \beta) f_2{}^{\rho_2 - 1}}$$

$\Longrightarrow$

$$f_2 = f_1 \left[\frac{b_2\beta}{b_1(1-\beta)}\right]^{\frac{1}{\rho_2 - 1}}$$

Plug in $L_\lambda = 0$, we have

$$f_1^* = \frac{L}{b_1 + b_2 A}, \qquad f_2^* = \frac{L \cdot A}{b_1 + b_2 A}$$

where $A = \left[\frac{b_2\beta}{b_1(1-\beta)}\right]^{\frac{1}{\rho_2 - 1)}}$.

D. Estimation Algorithm.

Let I denote the total number of bloggers, N is the number of previous iterations used for calculating the expected value for the current iteration.

1. At iteration r, the state-depend value $\tilde{V}$and heterogenous parameters $\boldsymbol{\Theta_i}$ in the past N iterations are $H^r = \{\{\boldsymbol{\Theta_i^m}, \tilde{V}^m(S_i^m; \boldsymbol{\Theta_i^m})\}_{i=1}^{I}\}_{m=r-N}^{r-1}$, where $\theta_i \in \boldsymbol{\Theta_i} = \{\alpha_{i1}, \alpha_{i2}, \boldsymbol{\rho_i}, \boldsymbol{\eta_i}, \boldsymbol{t_i}, \boldsymbol{\gamma_i}, \delta_i\}$.

2. Draw μ_θ^r, the population mean of θ_i, from the posterior distribution based on σ_θ^{r-1} and parameters estimated in the last iteration $\{\theta_i^{r-1}\}_{i=1}^{I}$, i.e., $\mu_\theta^r \sim N\left(\frac{\sum_{i=1}^{I}\theta_i^{r-1}}{I}, \sigma_\theta^{r-1}\right)$.

3. Draw σ_θ^r, the population variance of θ_i, from the posterior distribution based on the updated μ_θ^r and $\{\theta_i^{r-1}\}_{i=1}^{I}$, i.e., $\sigma_\theta^r \sim IG\left(\frac{I}{2}, \frac{\sum_{i=1}^{I}\left(\theta_i^{r-1} - \mu_\theta^r\right)^2}{2}\right)$.

4. Draw new parameters θ_i^r for each individual $i = 1, \dots, I$ from the posterior distribution $f_i(\theta_i^r | \mu_\theta^r, \sigma_\theta^r, {St_i}^o, {Br_i}^o, {P_i}^o) \propto \pi(\theta_i | \mu_\theta^r, \sigma_\theta^r) L\{{St_i}^o, {Br_i}^o, {P_i}^o | \theta_i^r\}$, where ${St_i}^o, {Br_i}^o, {P_i}^o$ are the observed choices of style, brand, and post.

We use Metropolis-Hastings algorithm to draw from the above posterior distribution.

(1). Draw candidate parameters θ_i^{*r} from a proposal density $q(\theta_i^{r-1}, \theta_i^{*r})$, essentially adding some perturbation to θ_i^{r-1}, for example, $\theta_i^{*r} \sim N(\theta_i^{r-1}, \varepsilon^2)$.

(2). Given θ_i^{*r}, compute the likelihood, i.e., $L\{{St_i}^o, {Br_i}^o, {P_i}^o | \theta_i^{*r}\}$. Computation of the likelihood with the observed continuous choice (${St_i}^o$) is traditionally done in two ways: first, discretizing the continuous choice space into countable discrete choices; second, numerical approximation using kernel smoothing (e.g., Yao et al., 2012; Liu et al., 2018). The downside of the first way is the loss of the continuous nature of the corresponding choice. The second approach reserves continuity but requires thousands of draws of random errors and solving for optimal choices, which could be computationally very costly, especially for the case without closed-form solutions.

Our approach borrows the spirits of both ways. On the one hand, we reserve the continuous nature of the style choices ${St_i}^o$, that is, we allow the choices to take any positive real numbers. On the other hand, to alleviate the computational burden, we also use some numerical approximation or smoothing, based on the 'discretizing' the number of likes which naturally take discrete values (i.e., integers).

Specifically, for a given target likes $\widehat{Like}$, there exists a unique choice $St^{*}{}_{i}$ (please see Appendix C for the optimal solution). The $\widehat{Like}$ can take any integer from 0 to the maximum of likes achieved across all the posted fashion looks $K = 1747$. The probability for the observed style choices $St_i{}^{o}$ that lead to $\widehat{Like} = \hat{g}(St_i{}^{o}, Br_i{}^{o}, S_{it})$ is therefore

$$Pr\{St_i{}^{o}, Br_i{}^{o}, P_i{}^{o} | S_{it}, \theta_i^{*r}\}$$

$$= \frac{\exp\left\{v^r\left(\widehat{Like}, Br_i{}^{o}, P_i{}^{o} | S_{it}, \theta_i^{*r}\right)\right\}}{\exp\left\{v^r(P_i{}^{o} = 0 | S_{it}, \theta_i^{*r})\right\} + \sum_{k=0}^{K}\sum_{n=1}^{35}\exp\left\{v^r\left(\mathrm{k}, Br_{i,n}, P_i{}^{o} = 1 | S_{it}, \theta_i^{*r}\right)\right\}}$$

where $v^r\left(\widehat{Like}, Br_i{}^{o}, P_i{}^{o} | S_{it}\right) = v^r\left(St_i{}^{o}, Br_i{}^{o}, P_i{}^{o} | S_{it}\right)$ if $\widehat{Like} = \hat{g}(St_i{}^{o}, Br_i{}^{o}, S_{it})$ is an integer. Otherwise, $v^r\left(\widehat{Like}, Br_i{}^{o}, P_i{}^{o} | S_{it}\right) = v^r\left(\widehat{Like\_nbr}, Br_i{}^{o}, P_i{}^{o} | S_{it}\right) = v^r\left(St^{*}{}_{i}, Br_i{}^{o}, P_i{}^{o} | S_{it}\right)$, where $\widehat{Like\_nbr}$ is the nearest integer neighbor of $\widehat{Like}$, and $St^{*}{}_{i}$ is the optimal choices to achieve the target $\widehat{Like\_nbr}$.

The choice specific value function $v^r(St_i{}^{o}, Br_i{}^{o}, P_i{}^{o} | S_{it})$ is the per-period utility plus the expected future value $EV^r(S_i)$, calculated with a weighted average of the past state-specific value $\{\{\tilde{V}^m(S_i^m; \boldsymbol{\Theta_i^m})\}_{i=1}^{I}\}_{m=r-N}^{r-1}$. The cumulative number of likes, denoted by s, as the stochastically evolving state, follow the state transition probability $T(St_i, Br_i, P_i | s, \hat{\sigma})$. We have the estimated standard deviation for the random error $\zeta_{it} \sim N\left(0, \hat{\sigma}^2\right)$, resulting from the transition regression.[38] With a step size of 1, the probability of getting y likes rather than $\widehat{Like}$ is therefore $1 \times \phi_{(0,\hat{\sigma})}(y - \widehat{Like})$, where $\phi_{(0,\hat{\sigma})}(\cdot)$ is the density function for normal distribution $N\left(0, \hat{\sigma}^2\right)$. So, we have

[38] An ordinal logistic regression was also tested for robustness check.

$EV^r(St_i, Br_i, P_i, S_i) = \sum_{m=r-N}^{r-1} \tilde{V}^m(S_i^m; \boldsymbol{\Theta_i^m}) \frac{K_{h\theta}(\theta_i^{*r} - \theta_i^m) K_{hS}(s_i^{m\prime} - s_i)}{\sum_{j=r-N}^{r-1} K_{h\theta}(\theta_i^{*r} - \theta_i^j) K_{hS}(s_i^{j\prime} - s_i)}.$

Then we can calculate the likelihood $L\{St_i{}^o, Br_i{}^o, P_i{}^o | \theta_i^{*r}\}$.

(3) Repeat the above to obtain the likelihood for the old parameter from the last iteration $L\{St_i{}^o, Br_i{}^o, P_i{}^o \big| \theta_i^{r-1}\}$.

(4) Having obtained the likelihood, we can determine whether to accept the candidate parameters $\hat{\theta}_i^{*r}$, with acceptance probability λ given by

$$\lambda = \min\left\{\frac{\pi(\theta_i^{*r} | \mu_\theta^r, \sigma_\theta^r) L\{St_i{}^o, Br_i{}^o, P_i{}^o | \theta_i^{*r}\} q\left(\theta_i^{*r}, \hat{\theta}_i^{r-1}\right)}{\pi\left(\theta_i^{r-1} | \mu_\theta^r, \sigma_\theta^r\right) L\{St_i{}^o, Br_i{}^o, P_i{}^o \big| \theta_i^{r-1}\} q(\theta_i^{r-1}, \theta_i^{*r})}, 1\right\}.$$

Repeat (1) ~ (4) for all individuals $i = 1, \ldots, I$.

5. Given the accepted parameters resulting from step 4, θ_i^r, we can update the value function for each individual, $\{\tilde{V}^r(S_i^r; \boldsymbol{\Theta_i^r})\}_{i=1}^I$.

Given the Type-1 extreme value distribution for the brand-and-style specific random error, the value takes the following form

$$\tilde{V}^r(S_i^r; \boldsymbol{\Theta_i^r}) = 0.577 + \log\left[\exp\{v^r(P_i{}^o = 0 | S_i, \theta_i^r)\} + \sum_{k=0}^{K} \sum_{n=1}^{35} \exp\left\{v^r\left(\mathrm{k}, Br_{i,n}, P_i{}^o = 1 | S_i, \theta_i^r\right)\right\}\right]$$

6. Proceed to the next iteration $r + 1$ and repeat the above steps until convergence.