

STRATEGIC ANALYSES OF USER GENERATED CONTENTS AND PLATFORMS

Submitted to the Tepper School of Business
in partial fulfillment for the requirements for the degree of
DOCTOR OF PHILOSOPHY in Industrial Administration

by

Baojun Jiang

Tepper School of Business
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213

April 2011

Committee Members:

Kannan Srinivasan (Chair)
Timothy P. Derdenger
Kinshuk Jerath
Baohong Sun

ACKNOWLEDGEMENT

I would like to thank my parents for all that they have taught me and for all their hard work of raising me and my two brothers. Without their unconditional love and support, I would never have reached where I am today.

I would like to thank my advisor, Kannan Srinivasan—the one person who made the most difference in my academic life—for his absolutely candid advice, guidance, push/encouragement, and support. I would also like to thank all other committee members—Tim Derdenger, Kinshuk Jerath, and Baohong Sun—for feedback on my research, and Nitin Mehta for so many direct, valuable questions and comments. My research has also benefited from the helpful suggestions from many other Carnegie Mellon faculty, Ph.D. students and alumni as well as marketing journal editors, associate editors, and anonymous reviewers. Any errors are mine and only mine.

I am very grateful for the fellowship and other support from the Tepper School of Business. I also thank the Informs Society for Marketing Science (ISMS) for its financial support of my dissertation research via its Inaugural ISMS Doctoral Dissertation Proposal Award.

And last but not least, I would like to thank my wife, Hui, for all the joy she has brought into my life since the first day we met and for all her hard work in the past few years, and my sons, Shang and Kai, for all their patience and understanding or lack thereof.

TABLE OF CONTENTS

1. Introduction	1
2. Impact of Online Consumer Reviews on Competitive Marketing Strategies	
2.1. Introduction	5
2.2. Model	10
2.3. Effects of Consumer Reviews	19
2.4. Alternative Sequential Decision Making and Implications	27
2.5. Strategic Manipulation of Reviews	29
2.6. Limitations	31
2.7. Conclusion	32
Appendix for Chapter 2	33
3. Firm Strategies in the “Mid Tail” of Platform-Based Retailing	
3.1. Introduction	39
3.2. Review of Relevant Literature	46
3.3. Model	48
3.4. Analysis	54
3.5. Effects of Consumer Reviews	63
3.6. Robustness to Alternative Assumptions	66
3.7. Conclusion	68
Appendix for Chapter 3	71
Technical Appendix for Chapter 3	79
4. Ad-Supported Platform in a Vertically Differentiated Market	
4.1. Introduction	86
4.2. Monopoly Market	90
4.3. Competitive Market	98
4.4. Conclusion	109
Appendix for Chapter 4	111
References	122

1. Introduction

The commercialization of the Internet began only in 1995 but has since penetrated almost every aspect of the society far beyond its initial, very limited research and academic use. The Internet has drastically changed both culture and commerce enabling new ways of instant communication and social or commercial interactions among a vast number of people and businesses without the old, overwhelming geographic constraints. My dissertation consists of three essays (chapters 2-4) that study the economic and strategic issues in the Internet technology-enabled markets.

Chapter 2 studies consumer product reviews—one of the earliest forms of online user generated contents. Consumer reviews have now become widely available on popular online retailer websites as well as on many third-party consumer information sharing websites. Such reviews play a significant role in consumer buying decisions as they can help consumers resolve or reduce uncertainties about product features and qualities before their purchases. Online reviews are becoming increasingly important for newly introduced products or products from less well-known companies. For example, when Motorola introduced Droid, its first Android smartphone, in late 2009 to compete with Apple's iPhone, consumers had generally given very positive reviews on Droid's key features such as its high resolution touch screen and camera, multitasking capabilities, open Android platform, and built-in Google Navigation. Given the positive reviews, how should Motorola adjust its pricing, promotion, and advertising strategies? For instance, should Motorola view these reviews as free advertising and thus lower its advertising? Are there conditions under which Motorola may actually increase its advertising expenditure in spite of its favorable reviews? How should Motorola adjust the level of its promotional premium—goods offered to consumers either free or at a low cost as an incentive to buy a product—in response to its favorable reviews? Should it increase its price since favorable reviews have increased consumers' willingness-to-pay? How should other smartphone makers respond? How do reviews affect firms' profits?

In Chapter 2, I provide a normative, game-theoretic model to study how consumer reviews influence firms' advertising, product premium, and pricing strategies. My analysis shows several

interesting findings. First, consumer reviews and firms' advertising should not affect their product premium strategies. Second, even though favorable reviews and advertising are *substitutes* with respect to increasing consumers' willingness to pay, the firm may actually consider them as *complements*, especially when advertising is expensive. Competitive firms' advertising responses to the availability of reviews are opposite: one firm will increase advertising whereas its competitor will decrease it. Third, the total industry advertising expenditure may increase even when both firms have positive reviews. Fourth, the effect of consumer reviews on the difference between firms' prices can be in the *opposite* direction to their effect on the relative separation between firms' perceived qualities. Lastly, because of the competitive responses in advertising, an improvement in a firm's reviews may hurt its profit and increase its competitor's profit.

In Chapter 3, I study an important phenomenon on online retail platforms (such as Amazon.com). While millions of products are sold on its retail platform, Amazon itself stocks and sells only a very small fraction of them. Most of these products are sold by third-party sellers, who pay Amazon a fee for each unit sold. Empirical evidence clearly suggests that Amazon tends to sell high-demand products and leave long-tail products for independent sellers to offer. I investigate how the platform owner, facing *ex ante* demand uncertainty, may strategically learn from these sellers' early sales which of the "mid-tail" products are worthwhile for its direct selling and which are best left for others to sell. The platform owner's "cherry-picking" of the successful products, however, gives an independent seller the incentive to mask any high demand by lowering his sales with a reduced service level (unobserved by the platform owner).

I analyze this strategic interaction between the platform owner and the independent seller using a game-theoretic model with two types of sellers—one with high demand and one with low demand. I show that it may not always be optimal for the platform owner to identify the seller's demand. Interestingly, the platform owner may be *worse off* by retaining its option to sell the independent seller's product whereas both types of sellers may *benefit* from the platform owner's threat of entry. The platform owner's entry

option may reduce the consumer surplus in the early period though it increases the consumer surplus in the later period. I also investigate how consumer reviews influence the market outcome.

In the last chapter, I study the Internet-enabled ad-supported licensing model—one that is becoming more popular in many software and application services markets. With the tremendous growth in network-enabled mobile computing, many large companies are striving to make their devices or software systems as a platform on which a huge number of consumer applications or services can be offered by third-party developers or providers. For example, in Apple’s App Store (for iPhone and iPad) and Google’s Android Market (for competing AndroidOS-based smartphones and tablet computers), we see hundreds of thousands of free ad-supported applications as well as ad-free, paid applications. And it is also common that the same application is offered in both a paid ad-free version and a free ad-supported version. Even in some hardware markets, companies are beginning to test the ad-supported model. For example, Amazon.com has just announced that it will offer a cheaper ad-supported version of Kindle, its bestselling electronic reader. All aspects of the Kindle are the same only that for the ad-supported model, special offers and sponsored screensavers will display on its screensaver and at the bottom of the home screen (without interrupting reading on it).¹ This may be Amazon’s first key step in turning its Kindle eReader into an advertising platform. The ad-supported model is clearly an important phenomenon in these new technology and Internet-enabled markets.

I study the adoption of the ad-supported model in these software applications markets. I show that, ignoring fixed costs, it is generally sub-optimal for a monopolist to offer only ad-free software. If the per-user advertising rate is high relative to consumers’ distaste for advertisements, the monopolist will offer only ad-supported software at a reduced price or for free; otherwise, it will offer both versions of its application. My analysis of a competitive vertically differentiated market, in which each firm adopts only one platform, shows that, unless one firm’s product is far inferior, both firms are better off if either firm adopts the ad-supported platform than if neither does. When both firms can potentially adopt multiple

¹ <http://www.usatoday.com/tech/products/2011-04-11-amazon-kindle-ads.htm>, April 11, 2011.

platforms, I find that, under very general conditions, the low quality firm offers only ad-free software whereas the high quality firm offers both ad-free and ad-supported software. More interestingly, I find that even if *neither* firm earns any positive advertising revenue or only one firm does, *both* firms can benefit from the availability of the ad-supported platform. My analysis suggests that, in a quality differentiated software market with intense price competition, firms may have incentives to adopt the ad-supported platform even if their advertising revenue does not cover the fixed cost required for that platform.

2. Impact of Online Consumer Reviews on Competitive Marketing Strategies

2.1. Introduction

Commensurate with the rapid growth of the Internet, online user generated content (UGC) has grown dramatically. Online consumer reviews, one of the earliest forms of UGC, have now become widely available on most popular retailer websites as well as on many third-party consumer information-sharing websites. Such reviews play a significant role in consumer buying decisions as they can help consumers resolve or reduce uncertainties about product features and qualities before their purchases. Business Week reports that 70% of Americans consult product reviews or consumer ratings before making a purchase (Ante 2009). An IPC Media study of UK female Internet users finds that 97% of all women research products online and 92% have bought products online (eMarketer 2009). Moreover, 92% of shoppers have more confidence in the product information they seek online than from other sources (Penn and Zalesne 2009). A survey of more than 2000 US online users shows that 97% of the consumers who made purchases based on online reviews said that they found the reviews to be accurate (comScore and The Kelsey Group 2007). The same survey reveals that consumers are willing to pay at least 20% more for service products with a 5-star rating than those with a 4-star rating. Consumer-generated reviews are becoming increasingly important for new products or products from less well-known companies. While it is evident how consumer reviews influence consumer buying decisions, it is not clear how consumer reviews may affect firms' marketing strategies. For example, when Motorola Inc. introduced its Droid smartphone in late 2009 to compete with Apple Inc.'s iPhone, consumers had generally given very positive reviews to the Droid's key features such as its high resolution touch screen, multitasking capabilities, open Android platform, and built-in Google Navigation.² Given the positive reviews, how should Motorola adjust its pricing, promotion, and advertising strategies? For example, should Motorola view these reviews as free advertising and thus lower its advertising expenditure? Are there conditions under which Motorola may actually increase its advertising in spite of its favorable reviews? How should

² See, for example, <http://www.wirelessweek.com/Reviews/2009/11/Review--Droid-vs--the-iPhone-3GS/> and <http://reviews.cnet.com/motorola-droid-review/>, last accessed in June 2010.

Motorola adjust the level of its promotional premium—goods offered to consumers either free or at a low cost as an incentive to buy a product—in response to its favorable reviews? Should it increase its price since favorable reviews have increased consumers’ willingness-to-pay? How should other smartphone makers respond? How do reviews affect firms’ profits?

Our research focuses on the key question: How does UGC such as consumer product reviews affect firms’ advertising, product premium, and pricing strategies? Note that consumer reviews, advertising, and product premium can all influence a consumer’s purchase decision (i.e., her willingness-to-pay for the product), but they have very different cost implications to the firms. Consumer reviews are typically costless to the firms, while advertising represents a fixed cost and product premium imposes a variable cost (representing any such promotional tool). One of our main goals is to investigate whether the substitutability among these factors from the consumer’s perspective (in terms of changing her product valuation) necessarily leads to substitutions of these factors by the firms.

We provide a normative, game-theoretic model to address the following specific research questions. Do firms consider their positive or favorable reviews as a substitute for their advertising or product premium? When a firm receives favorable reviews relative to its competitor, how should it optimally adjust its advertising, product premium, and price? How will the competitor respond? How are firm profits affected by consumer reviews? How do consumer reviews affect the total advertising expenditure and the price separation in the industry? Does the firm with favorable reviews have less incentive to strategically manipulate its reviews?

Our research contributes to the growing literature on UGC, in particular on consumer reviews as a form of word-of-mouth (WOM). Prior empirical research has shown the association between online reviews/WOM and sales in various product categories, though conflicting evidence exists as to what measures of reviews/WOM influence sales. Godes and Mayzlin (2004) show that the dispersion of conversations about TV shows across online consumer communities is strongly associated with the rating (i.e. popularity) of these shows. Liu (2006) finds that the volume of WOM rather than the percentages of positive and negative messages explains box office revenues in the movie industry. Dellarocas et al.

(2007) find that both the volume and the average rating of reviews are statistically significant in predicting future movie revenues. Duan et al. (2008) identify a positive feedback loop between online WOM and box office revenue where the box office revenue affects the online WOM volume, which in turn affects the box office revenue. Chen et al. (2006) find that for books, higher ratings are associated with higher sales and that highly informative reviews (i.e., those with a high proportion of helpful votes) strengthen this effect with additional sales. The variance in ratings may also influence the role of product ratings for the consumers and the firms (e.g., Clemons et al. 2006; Martin, Barron and Norton 2007; Sun 2009). Chevalier and Mayzlin (2006) find that negative (one-star) reviews have greater impact on book sales than positive (five-star) reviews and that consumers may read and respond to written text reviews rather than merely the average star rating. Forman et al. (2008) provide evidence that consumers find reviews containing identity-descriptive information more helpful and that the prevalence of reviewer disclosure of identity information is positively associated with subsequent product sales. With data from the video game industry, Zhu and Zhang (2010) find that even in the same product category, the impact of reviews can differ across products. While a number of empirical papers show a significant association between reviews and sales, thus far there has been limited effort to address how firms' optimal marketing strategies are affected by reviews. Chen and Xie (2008) study how consumer reviews influence a monopolist's information content strategy in a market with expert and novice users. They show that seller-created product information and consumer-created product information (i.e., consumer reviews) can be complements or substitutes to the monopoly seller. That is, when reviews become available, the monopolist may increase or decrease the release of its product attribute information depending on the size of the expert user segment. Jiang and Wang (2008) show that the effect of consumer reviews on a firm's profit depends on the informativeness of the reviews, the quantifiability of the product attributes, and the competitive environment. Kuksov and Xie (2010) find that a monopolist should adapt its strategy if consumers do not observe the price history from customer ratings and that its optimal pricing and "frills" strategies depend on the market growth rate.

In an important and pioneering research effort, Chen and Xie (2005) study third-party product reviews and find that firms should respond to product reviews by adjusting their advertising expenditures but not their prices. This result critically depends on the discrete two-level consumer valuations in their model. We show that with continuous consumer valuations, firms adjust both advertising expenditures and prices in response to reviews. In addition, we allow for advertising and product reviews to directly change each consumer's *ex ante* product valuation; they model advertising and reviews as affecting the aggregate fraction of consumers who hold correct beliefs about product valuations. Finally, we examine the effect of product reviews on other marketing mix variables such as product premium. As noted earlier, product premium represents any promotional activity that results in a variable cost and thus our result generalizes to any such promotional tool. We show that firms should not adjust their product premium strategies as reviews for their core product change. In our model, the competing firms differ in terms of the effectiveness of advertising, the cost of premium and baseline perceived quality levels. Therefore, our analysis leads to substantive insights when there is a significant difference between the firms on these strategic levers. We are able to identify not-so-obvious insights. For example, when there is a difference in the firms' perceived qualities before reviews are available, and reviews make that difference even larger, one might expect the price difference between the firms' products to widen as well. Interestingly, we show that the price difference may actually narrow.

With a comprehensive framework adopted in this chapter, we are able to shed insight on the incentives for firms to manipulate their reviews. Mayzlin (2006) finds that a firm with inferior products will spend more resources on promotional chat activities than a firm with superior products. Dellarocas (2006) shows that a higher quality firm may also have higher incentives to manipulate ratings and hence strategic manipulation of reviews can increase the informativeness of online forums. Li and Hitt (2008) show that the self-selection bias in early reviews may reduce consumer surplus if later consumers fail to correct that bias. Kuksov and Xie (2010) show how a monopolist should manipulate its rating through pricing or offering of frills to early consumers. Our research complements these works; we find that, in a

competitive market, a firm is more likely to complement its favorable reviews with additional manipulation when the cost of manipulation is high than when it is low.

Our research complements the aforementioned research by considering the interactions of consumer reviews with firms' competitive marketing strategies. We develop a normative, game-theoretic framework that incorporates consumer product reviews into firms' pricing, advertising and premium strategies to examine the research questions which have not been fully addressed in the literature. Our research shows several interesting findings. First, we find that consumer reviews and firms' advertising do not affect firms' product premium strategies. Second, we show that, even though a firm's favorable consumer reviews and advertising are substitutes with respect to increasing the consumer's willingness to pay, the firm may actually consider them as complements, especially when advertising is inefficient or expensive. The complementarities between the firm's favorable reviews and its advertising will strengthen if its favorable reviews also make its advertising more effective. In the case of Motorola's Droid, Motorola and Verizon (the wireless carrier) complemented and supported Droid's favorable reviews with a huge advertising budget, which helped Motorola exceed Apple in brand loyalty among adult men.³ Our analysis shows that competitive firms' advertising responses to the availability of reviews are opposite, i.e., one firm will increase advertising whereas the other firm will decrease it. Third, the total industry advertising expenditure may increase even when both firms receive positive reviews that increase consumers' product valuations. Fourth, counter-intuitively, the effect of consumer reviews on firms' price separation can be in the *opposite* direction to their effect on the relative separation between firms' perceived qualities. Fifth, surprisingly, an improvement in a firm's reviews may hurt its own profit and increase its competitor's profit. Sixth, the strength of the consumer's taste preference influences the impact of consumer reviews on advertising and firms' best advertising responses to a stronger consumer taste preference are qualitatively different. Lastly, a firm is more likely to complement its favorable

³ http://www.appleinsider.com/articles/09/11/24/motorola_passes_apple_in_brand_loyalty_among_men_study.html

reviews with additional manipulation of its reviews when the cost of manipulation is high than when it is low.

The rest of this chapter is organized as follows. In Section 2.2., we present our model framework and solve for the subgame perfect Nash equilibrium. In Section 2.3., we analyze the impact of consumer reviews on firms' marketing strategies. In Section 2.4., we study an alternative dynamic game as well as a static game to check for the robustness of our results. In Section 2.5., we extend our model to study firms' incentives to manipulate their reviews. We discuss the limitations of our model in Section 2.6. and conclude the chapter in Section 2.7.

2.2. Model

We study how consumer reviews influence firms' marketing strategies in a competitive duopoly market with differentiated products. We model two distinct time periods. In the first period, no consumer reviews are available and consumers make their purchase decisions based on each firm's first-period price, advertising, and product premium offering. In the second period, consumer reviews become available (after first-period consumers rate and review the products); the second-period consumers will make their purchases based on each firm's second-period price, advertising and premium offering as well as the consumer review information. We adopt the standard Hotelling model setting. Each firm offers one product—firm i offers product i and has a marginal cost of c_i . Without loss of generality, we assume firm 1's product is located at zero and firm 2's product at one. New consumers arrive in each period. We use x to represent the “location” or the horizontal preference of a consumer. Consumers are heterogeneous with respect to their preferences and are uniformly distributed on the line segment between zero and one: $x \sim \text{uniform}[0, 1]$. The consumer's disutility from non-exactly matched preference is td , where d is the distance between the consumer's location (x) and the product's location, and t represents the strength of consumers' preferences. Each consumer will buy at most one product. Without loss of generality, we normalize the total number of consumers in each period to be one. Our setup represents a scenario in which firms have already developed their products through their research and development efforts, and

now the question is how they should best adapt their marketing strategies to the availability of consumer reviews.

It is well established in the literature that product reviews can influence consumers' valuations of the product. We accommodate such effects in the following way. Let $\beta_i V_i$ be a perfectly matched ($d = 0$) consumer's valuation of product i , where the positive coefficient β_i is used to capture how reviews influence the consumer's valuation. We normalize β_i such that V_i represents an ideal consumer's *ex ante* valuation for product i without any consumer review information, i.e. $\beta_i = 1$ in the first period when reviews are not yet available. The second-period consumers will have access to consumer reviews; with consumer review information, the consumers' perception of quality and hence their product valuations may change. We use changes in β_i to capture how consumer reviews affect the consumers' willingness to pay in the second period: $\beta_i > 1$ if the reviews for product i are positive and $\beta_i < 1$ if the reviews are negative.⁴ Such changes in the consumers' valuation are consistent with the existing literature. Consumers' *ex ante* expected valuation of the product may change when they receive more information about the product. For example, a firm may proactively control the amount of product information released to consumers in an attempt to influence their valuations by changing their *ex ante* probability of finding the product to match their tastes (Chen and Xie 2008; Lewis and Sappington 1994). In the current context, consumer review information may help consumers resolve uncertainties about product qualities and thereby affect their product valuations. In addition, reviews may be considered as free advertising that fosters brand preferences, and hence will increase (or decrease if the reviews are negative) the consumer's willingness to pay. We take this reduced form approach to modeling consumer reviews since our goal is to study the economic impact of reviews on firms' marketing strategies rather than to study the evolutionary mechanism of reviews per se.

In each period, each firm i has three decision variables—price (p_i), the level of advertising (a_i) and the level of product premium (e_i). As in Adams and Yellen (1977), we assume that advertising by

⁴ Note that instead of using the form $\beta_i V_i$ to capture the effect of reviews on consumers' valuation, we can equivalently use $V_i + R_i$, where R_i represents the change in consumers' valuation caused by reviews. Our results are the same with the alternative formulation.

firm i can increase consumers' willingness to pay for its product by an amount a_i . This assumption is also consistent with the persuasive view of advertising—advertising alters consumers' preferences and creates brand loyalty—and the complementary view of advertising—advertising (when properly carried out) provides additional utility to consumers such as creating a feeling of greater social prestige (Bagwell 2007). We use a general convex function $g_i(a_i)$ to denote firm i 's fixed cost required for its advertising level a_i . Further, a firm's advertising level a_i in the first period may have also some residual effect (δa_i) on the second-period consumers' valuations. Product premium is a promotional tool. We abstract the premium to be some good of value e_i that firm i offers to consumers together with its product. That is, when firm i sets its product premium level at e_i , consumers' valuation for its product will increase by an amount e_i . For example, if a consumer who buys a laptop computer also receives a free carrying bag, the consumer's valuation for the carrying bag corresponds to e_i . In this case, different e_i levels may correspond to the quality of the bag (e.g., whether it is made of nylon or leather). For a premium level of e_i , firm i needs to incur an extra marginal cost (in addition to c_i), which is denoted by a continuously differentiable function $f_i(e_i)$. Clearly, $f_i(0) = 0$ since firms incur no premium costs if they do not offer any premium. We also assume that firms have a bounded efficiency at offering premium; stated formally and more restrictively, there exists some $\varepsilon > 0$ such that $e_i - f_i(e_i) < 0$ for all $e_i > \varepsilon$. In addition, since we study a competitive market rather than a monopoly, we implicitly assume that firms' efficiencies in advertising or offering premium are not so vastly different that one firm is able to profitably force its competitor out of the market through its advertising or premium offering. From the firm's perspective, the salient difference between advertising and premium offering is that advertising entails a fixed cost whereas product premium has a variable cost.

In summary, given firms' advertising (a_i), product premium (e_i), and price (p_i) decisions, a first-period consumer of type x will derive a net utility of $U_i = \beta_i^{(1)}V_i + a_i^{(1)} + e_i^{(1)} - td - p_i^{(1)}$ from product i , where the parenthesized superscripts indicate the time period and where $d = x$ for $i = 1$ and $d = 1 - x$

for $i = 2$.⁵ In the second period, a consumer of type x derives a net utility of $U_i = \beta_i^{(2)}V_i + \delta a_i^{(1)} + a_i^{(2)} + e_i^{(2)} - td - p_i^{(2)}$ from product i . Note that $\delta a_i^{(1)}$ is the residual effect from firm i 's first period advertising. For analytical simplicity, we will analyze the $\delta = 0$ case; however, our core results remain qualitatively the same even with positive δ . This simplification allows us to combine the analysis for the two periods into one set of notations. Since new consumers come in each period, with $\delta = 0$, each firm will effectively make its marketing decisions to maximize its profit in each period separately. Thus, for expositional conciseness, we will, from now on, drop the parenthesized superscripts that indicate the time period. We assume that consumers' outside option has a utility of zero; hence, consumers will buy the product that yields a higher non-negative utility.

Note that, if the product valuations $(\beta_i V_i + a_i + e_i)$ are too low, both firms will be localized monopolies and the market will not be covered. Further, if the difference in the two products' valuations is too large, one firm will profitably squeeze the other out of the market and become a monopoly. Since our goal is to examine how consumer reviews affect firms' marketing decisions in the more interesting case of a competitive market, we will implicitly assume that the product valuations are not too low and that the difference in the two products' valuations is not too large. This assumption is equivalent to assuming that the strength (t) of consumers' horizontal preference is not too small and not too large. We will be more explicit about this assumption later.

Our specification of the direct effect of consumer reviews on the consumer's utility is simple and straightforward. The distinction between advertising and consumer reviews is subtle but not substantially different. On the demand side, product premium, advertising, and positive consumer reviews can all increase consumers' willingness to pay. Will firms from the supply side then consider these three factors to be substitutes as well? As can be seen from the consumer utility expression U_i , we have specified that

⁵ Though the *levels* of advertising and premium are added linearly to the consumer's utility, we actually model these factors in nonlinear ways (through the nonlinear cost functions associated with these levels), e.g., there is a diminishing return of advertising. An alternative and equivalent formulation is to use the advertising expenditure (S_i) and the marginal cost of premium (C_i) as firms' decision variables rather than a_i and e_i ; the consumer's utility will then be nonlinear in S_i and C_i : $U_i = \beta_i V_i + g^{-1}(S_i) + f^{-1}(C_i) - td - p_i$.

these factors can be perfect substitutes from the consumers' perspective. From the firm's perspective, however, product premium imposes a variable cost, advertising a fixed cost, and consumer reviews no cost (barring firms' strategic manipulations of reviews that we will study later). Intuitively, our current specification should make it more likely that firms will, to some extent, substitute among positive reviews, its level of premium, and advertising. Yet, we show later that even with the assumption of perfect substitutability from the consumer side, firms will not change its premium offering in response to reviews and they may consider positive consumer reviews and advertising as *complements* rather than *substitutes*.

Firms have three decisions to make (in each period)—advertising levels, product premium levels, and prices. Clearly, prices are the easiest to change among the three, so we assume that pricing decisions are made after advertising and product premium decisions. We also note that, in practice, fixed cost decisions tend to be made earlier than variable cost decisions; thus we assume that advertising (a fixed cost) is determined before product premium (a variable cost). Therefore, our game has three stages. In the first stage, firms simultaneously choose their respective advertising levels. In the second stage, they simultaneously choose their product premium levels. In the third stage, firms simultaneously set prices and consumers subsequently make purchase decisions. Later, we show that the equilibrium outcome is the same even if the firms make both advertising and product premium decisions together (i.e., the second and the third stage are combined). We will also show that our results remain qualitatively the same even if the game is static (i.e., both firms simultaneously set their respective price, advertising and product premium levels). We use the standard backwards induction technique to solve for the pure-strategy subgame perfect Nash equilibrium.

2.2.1. Competitive Pricing Decisions

Let x_{in} be the consumer who is indifferent between the two products (i.e., $U_1 = U_2$):

$$\beta_1 V_1 - t x_{in} - p_1 + a_1 + e_1 = \beta_2 V_2 - t(1 - x_{in}) - p_2 + a_2 + e_2 \quad . \quad \text{We easily find that } x_{in} = \frac{\beta_1 V_1 + a_1 + e_1 - \beta_2 V_2 - a_2 - e_2 + t - p_1 + p_2}{2t} .$$

Consumers with $x \leq x_{in}$ prefer to buy product 1 and those with $x > x_{in}$ prefer product 2. Firms' market shares (m_i) are x_{in} and $1 - x_{in}$, respectively; their profits are as follows.

$$\Pi_1(p_1, p_2, e_1, e_2, a_1, a_2) = [p_1 - c_1 - f_1(e_1)]x_{in} - g_1(a_1) \quad (1)$$

$$\Pi_2(p_1, p_2, e_1, e_2, a_1, a_2) = [p_2 - c_2 - f_2(e_2)](1 - x_{in}) - g_2(a_2) \quad (2)$$

The equilibrium outcome a_i^* , e_i^* , and p_i^* is solved via backwards induction. In the third stage, firms simultaneously choose prices given their advertising and product premium levels to maximize their respective profits. Note that the profit functions are inverted quadratic in prices. Firms' optimal (equilibrium) prices are obtained by simultaneously solving the first order conditions: $\frac{\partial \Pi_1}{\partial p_1} = 0$ and $\frac{\partial \Pi_2}{\partial p_2} = 0$.

$$p_1^*(e_1, e_2, a_1, a_2) = \frac{\beta_1 V_1 + a_1 + e_1 - \beta_2 V_2 - a_2 - e_2 + 3t + 2c_1 + c_2 + 2f_1(e_1) + f_2(e_2)}{3} \quad (3)$$

$$p_2^*(e_1, e_2, a_1, a_2) = \frac{\beta_2 V_2 + a_2 + e_2 - \beta_1 V_1 - a_1 - e_1 + 3t + c_1 + 2c_2 + f_1(e_1) + 2f_2(e_2)}{3} \quad (4)$$

We now digress to discuss our implicit assumption of a competitive (rather than monopoly) equilibrium in which both firms have a positive market share. For any firm i to profitably sell its product, its price must cover its marginal cost: $p_i^*(e_1, e_2, a_1, a_2) > c_i + f_i(e_i)$, $i \in \{1, 2\}$; otherwise, the firm will prefer not selling and have a zero market share. The necessary condition for a competitive pricing equilibrium straightforwardly simplifies to

$$|\beta_i V_i + a_i + e_i - \beta_j V_j - a_j - e_j - c_i - f_i(e_i) + c_j + f_j(e_j)| < 3t, \quad (A1)$$

where $i, j \in \{1, 2\}$ and $i \neq j$.⁶ Note that $\beta_i V_i + a_i + e_i - c_i - f_i(e_i)$ is firm i 's maximum feasible profit margin—the highest willing-to-pay consumer's valuation subtracted by firm i 's marginal cost—and that this essentially represents how competitive firm i can be. Thus, the condition (A1) is quite intuitive; it simply states that the difference in firms' feasible profit margins must be below a threshold; otherwise,

⁶ In this paper, when the subscripts i and j appear in the same expression, it is always assumed that $i \neq j$.

there may not be a pure-strategy pricing equilibrium, or the more competitive firm will effectively be a monopoly.⁷

Substituting (3) and (4) into (1) and (2), we simplify the firms' profits to

$$\Pi_1^*(e_1, e_2, a_1, a_2) = \frac{[3t + \beta_1 V_1 + a_1 + e_1 - \beta_2 V_2 - a_2 - e_2 - c_1 + c_2 - f_1(e_1) + f_2(e_2)]^2}{18t} - g_1(a_1) \quad (5)$$

$$\Pi_2^*(e_1, e_2, a_1, a_2) = \frac{[3t - \beta_1 V_1 - a_1 - e_1 + \beta_2 V_2 + a_2 + e_2 + c_1 - c_2 + f_1(e_1) - f_2(e_2)]^2}{18t} - g_2(a_2) \quad (6).$$

Firms' market shares are given by $m_1 = x_{in}^* = \frac{3t + \beta_1 V_1 + a_1 + e_1 - \beta_2 V_2 - a_2 - e_2 - c_1 + c_2 - f_1(e_1) + f_2(e_2)}{6t}$ and $m_2 = 1 - m_1 = \frac{3t + \beta_2 V_2 + a_2 + e_2 - \beta_1 V_1 - a_1 - e_1 + c_1 - c_2 + f_1(e_1) - f_2(e_2)}{6t}$, respectively.

2.2.2. Competitive Product Premium Decisions

In the second stage of the game, firms simultaneously choose their respective product premium levels (e_i) taking their advertising levels (a_i) as given. Proposition 1 shows that firms' equilibrium product premium levels are given by

$$e_i^* = \operatorname{argmax}_{e_i \geq 0} \{e_i - f_i(e_i)\}, i \in \{1, 2\}. \quad (7)$$

Note that since firms' equilibrium premium levels do not depend on their advertising, equation (7) will also form part of the equilibrium outcome of the complete game. We will keep e_i^* implicit throughout the chapter, but one can easily solve for e_i^* for any specific function $f_i(e_i)$; for example, if $f_i(e_i) = b_i e_i^2$, then $e_i^* = \frac{1}{2b_i}$.

Proposition 1: Product Premium Equilibrium

Firms' equilibrium levels of product premium do not depend on their advertising levels; consumer product reviews do not affect firms' optimal product premium decisions. Mathematically, $e_i^ = \operatorname{argmax}_{e_i \geq 0} \{e_i - f_i(e_i)\}, i \in \{1, 2\}$.*

⁷ Condition (A1) implies that the "transportation cost" (t) is not too small, i.e., the two firms are not too closely located. D'Aspremont et al. (1979) demonstrate that if two firms are closely located, there is no pure strategy equilibrium in the price subgame (with zero marginal costs). In our model, since firms are asymmetric with positive marginal costs, there are two possibilities when condition (A1) is violated—non-existence of a pure strategy pricing equilibrium or a monopoly price equilibrium (with the more competitive firm effectively setting a monopoly price and the other firm pricing at its marginal cost but getting zero market share).

We find that firms' equilibrium product premium levels are independent of their advertising levels. Each firm's best product premium strategy is to pick the level of premium that achieves the highest contribution to its own feasible profit margin. Stated differently, a firm's optimal variable-cost (product premium) decision does not depend on its fixed-cost decision (advertising). The intuition is that firms will adjust their pricing strategies to compensate for their product premium levels. Hence a firm's optimal product premium strategy is simply to obtain the highest valuation minus cost margin, $e_i - f_i(e_i)$, given its own cost structure, because if a firm does not choose such a premium level, it will be at a disadvantage to its competitor when setting its price. Further, though reviews influence consumers' perceived product qualities and their willingness to pay, they do not affect firms' product premium decisions in a competitive market.

The resulting profits for the current subgame are given by

$$\Pi_1^*(a_1, a_2) = \frac{[3t + \beta_1 V_1 + a_1 + e_1^* - \beta_2 V_2 - a_2 - e_2^* - c_1 + c_2 - f_1(e_1^*) + f_2(e_2^*)]^2}{18t} - g_1(a_1) \quad (8)$$

$$\Pi_2^*(a_1, a_2) = \frac{[3t - \beta_1 V_1 - a_1 - e_1^* + \beta_2 V_2 + a_2 + e_2^* + c_1 - c_2 + f_1(e_1^*) - f_2(e_2^*)]^2}{18t} - g_2(a_2) \quad (9).$$

2.2.3. Competitive Advertising Decisions

We now analyze the first stage of the game, in which firms simultaneously choose their advertising levels (a_i) to maximize their profits (8) and (9), respectively. To facilitate closed-form analytical solutions, we assume that firms' advertising costs are quadratic (e.g., Tirole 1988, Bagwell 2005): $g_i(a_i) = k_i a_i^2$, $i \in \{1, 2\}$. For the existence of a competitive equilibrium, advertising cannot be too "effective" at increasing consumers' valuations (i.e., k_i cannot be too close to zero). In particular, we assume $k_i > \frac{1}{18t}$, $i \in \{1, 2\}$, which implies that firms' profits (8) and (9) are *inverted* quadratic functions in terms of their advertising levels.

The first order conditions are

$$\frac{\partial \Pi_1^*(a_1, a_2)}{\partial a_1} = \frac{3t + \beta_1 V_1 + a_1 + e_1^* - \beta_2 V_2 - a_2 - e_2^* - c_1 + c_2 - f_1(e_1^*) + f_2(e_2^*)}{9t} - 2k_1 a_1 = 0 \quad (10)$$

$$\frac{\partial \Pi_2^*(a_1, a_2)}{\partial a_2} = \frac{3t - \beta_1 V_1 - a_1 - e_1^* + \beta_2 V_2 + a_2 + e_2^* + c_1 - c_2 + f_1(e_1^*) - f_2(e_2^*)}{9t} - 2k_2 a_2 = 0. \quad (11)$$

We can obtain the equilibrium advertising levels by simultaneously solving (10) and (11). Note that for assumption (A1) to hold, firms' advertising efficiencies cannot be too drastically different, because otherwise the efficient firm will be able to profitably use advertising to squeeze out its advertising-inefficient competitor by running a high enough level of advertising. So, both firms need to be comparably efficient at advertising. This consideration gives us two parameter regions of interest:

- (i) inefficient advertising market: $k_i > k_i^{(c)}$ for $i \in \{1, 2\}$,
- (ii) efficient advertising market: $\frac{1}{18t} < k_i < k_i^{(c)}$ for $i \in \{1, 2\}$, where for expositional conciseness, we have defined two constant expressions:

$$k_i^{(c)} \equiv \frac{1}{3[3t - \beta_i V_i + \beta_j V_j - c_i - c_j - e_i^* + f_i(e_i^*) + e_j^* - f_j(e_j^*)]} \text{ for } i, j \in \{1, 2\}.$$

The competitive equilibrium outcome as follows is obtained from interior solutions ($a_i^* > 0$ for $i = 1, 2$).⁸ Simultaneously solving (10) and (11) yields the equilibrium a_i^* ; the overall equilibrium outcome is as follows.⁹

$$a_i^* = \frac{3k_j[3t + \beta_i V_i - \beta_j V_j - c_i + c_j + e_i^* - f_i(e_i^*) - e_j^* + f_j(e_j^*)] - 1}{3(18k_i k_j t - k_i - k_j)} \quad (12)$$

$$p_i^* = \frac{2k_i t [3k_j(3t + \beta_i V_i - \beta_j V_j - c_i + c_j + e_i^* - f_i(e_i^*) - e_j^* + f_j(e_j^*)) - 1]}{18k_i k_j t - k_i - k_j} + c_i + f_i(e_i^*) \quad (13)$$

$$\Pi_i^* = \frac{k_i(18k_i t - 1)[3k_j(3t + \beta_i V_i - \beta_j V_j - c_i + c_j + e_i^* - f_i(e_i^*) - e_j^* + f_j(e_j^*)) - 1]^2}{9(18k_i k_j t - k_i - k_j)^2}, \quad (14)$$

where the equilibrium product premium levels are given by (7). In the Appendix of this chapter, we show that when firms have comparable advertising efficiencies, both will advertise and earn a positive profit in the subgame perfect Nash equilibrium. Intuitively, if neither firm is much more efficient at advertising

⁸ Boundary solutions (i.e., $a_i^* = 0$ for either firm) imply that the firm with zero advertising will get a zero market share at the pricing stage. The existence of interior solutions (rather than boundary solutions) results from our implicit assumption that firms' abilities to advertise are not drastically different, which is necessary for a *competitive* (rather than *monopoly*) equilibrium market outcome.

⁹ In the interest of space, we exclude from the paper any straightforward algebraic manipulations such as solving simultaneous linear equations, and will provide only the results of such manipulations.

than the other, then neither will be able out-advertise the other and both will advertise and make a profit at the unique equilibrium.

2.2.4. Pre-Review Period as a Benchmark

The pre-review, first period serves as an interesting and useful benchmark for the post-review, second period. When a new product is first offered in the market, consumer reviews are not available since no consumer has yet bought the product. As pointed out earlier, assuming *new* consumers arrive in each period, our analysis is the same as before and we only need to set $\beta_1 = 1$ and $\beta_2 = 1$ to indicate the lack of consumer review information. We use the superscript “*o*” to indicate the absence of consumer reviews. The equilibrium premium levels are still given by (7), and the equilibrium advertising levels, prices, and profits are as follows.

$$a_i^{*(o)} = \frac{3k_j[3t+V_i-V_j-c_i+c_j+e_i^*-f_i(e_i^*)-e_j^*+f_j(e_j^*)]-1}{3(18k_ik_jt-k_i-k_j)} \quad (15)$$

$$p_i^{*(o)} = \frac{2k_i[3k_j(3t+V_i-V_j-c_i+c_j+e_i^*-f_i(e_i^*)-e_j^*+f_j(e_j^*))]-1}{18k_ik_jt-k_i-k_j} + c_i + f_i(e_i^*) \quad (16)$$

$$\Pi_i^{*(o)} = \frac{k_i(18k_ik_jt-1)[3k_j(3t+V_i-V_j-c_i+c_j+e_i^*-f_i(e_i^*)-e_j^*+f_j(e_j^*))]-1]^2}{9(18k_ik_jt-k_i-k_j)^2} \quad (17)$$

2.3. Effects of Consumer Reviews

From the consumer’s perspective, a firm’s product premium (e_i) and advertising (a_i) can increase the consumer’s willingness to pay for its product. Positive consumer reviews (i.e., $\beta_i > 1$) can also increase the consumer’s product valuation; hence, on the demand side (as can be seen from the consumer surplus expression U_i), product premium, advertising, and positive consumer reviews are perfect substitutes. Will firms then consider these three factors to be substitutes as well? On the supply side (from the firm’s perspective), product premium imposes a variable cost, advertising a fixed cost, and consumer reviews no cost. According to Proposition 1, competitive firms do not consider advertising and product premium as substitutes; in other words, a firm’s best variable-cost strategy is independent of its best fixed-cost strategy. The same proposition also reveals that consumer reviews do not influence firms’

optimal product premium strategies. Intuitively, this is because each firm must still do its best at value creation through product premium (i.e., maximize $e_i - f_i(e_i)$), otherwise it will be at a disadvantage when pricing its product against the competitor.

Now we examine how consumer reviews may influence firms' advertising and pricing strategies. Should firms reduce their brand advertising to save costs when they receive positive reviews ($\beta_i > 1$)? Should they increase advertising to compensate for their negative reviews?

By examining the functional forms of the equilibrium outcome, we see that consumer reviews influence equilibrium advertising, pricing, or profits only if they differentially alter the consumer's valuations for the firms' products. If reviews change consumers' valuations for both firms' products by the same amount (i.e., $\Delta\beta_i V_i = \Delta\beta_j V_j$), firms' equilibrium strategies (prices, advertising and product premium levels) will remain the same. In particular, because of competition, firms will not be able to profitably raise their prices even though their positive reviews have increased consumers' willingness-to-pay for their products albeit by an equal amount. Thus, when we examine the effects of consumer reviews, we should look at the extent to which reviews or changes in reviews influence the consumer's valuations for the two products differently. Note that $\frac{\partial a_i^*}{\partial(\beta_i V_i - \beta_j V_j)} = \frac{1}{V_i} \frac{\partial a_i^*}{\partial \beta_i}$, $\frac{\partial p_i^*}{\partial(\beta_i V_i - \beta_j V_j)} = \frac{1}{V_i} \frac{\partial p_i^*}{\partial \beta_i}$, and $\frac{\partial \pi_i^*}{\partial(\beta_i V_i - \beta_j V_j)} = \frac{1}{V_i} \frac{\partial \pi_i^*}{\partial \beta_i}$. Hence, when we examine comparative statics, the following two variations yield the same direction of change in the equilibrium variables:

- (i) an increase (decrease) in β_i given β_j ,
- (ii) an increase (decrease) in $\beta_i V_i - \beta_j V_j$.

Firm i has a favorable change in reviews if $\beta_i V_i - \beta_j V_j$ increases. Firm i is said to have more favorable reviews than firm j if $(\beta_i V_i - V_i) - (\beta_j V_j - V_j) > 0$ or equivalently if $(\beta_i V_i - \beta_j V_j) - (V_i - V_j) > 0$. Table 2.1 shows comparative statics in all parameter regions of interest. Proposition 2 sheds light on the effect of consumer reviews on firms' advertising and pricing strategies.

Table 2.1: Key Comparative Statics Regarding Consumer Reviews

Parameter Regions	Advertising	Price	Profit
Efficient Advertising: $\frac{1}{18t} < k_i < k_i^{(c)}$ for $i \in \{1, 2\}$	$\frac{\partial a_i^*}{\partial \beta_i} < 0$, $\frac{\partial a_i^*}{\partial \beta_j} > 0$	$\frac{\partial p_i^*}{\partial \beta_i} < 0$, $\frac{\partial p_i^*}{\partial \beta_j} > 0$	$\frac{\partial \pi_i^*}{\partial \beta_i} < 0$, $\frac{\partial \pi_j^*}{\partial \beta_i} > 0$
Inefficient Advertising: $k_i > k_i^{(c)}$ for $i \in \{1, 2\}$	$\frac{\partial a_i^*}{\partial \beta_i} > 0$, $\frac{\partial a_i^*}{\partial \beta_j} < 0$	$\frac{\partial p_i^*}{\partial \beta_i} > 0$, $\frac{\partial p_i^*}{\partial \beta_j} < 0$	$\frac{\partial \pi_i^*}{\partial \beta_i} > 0$, $\frac{\partial \pi_j^*}{\partial \beta_i} < 0$

Proposition 2: Effects of Consumer Reviews on Firm Prices and Advertising

When advertising costs are relatively high, firms consider their own favorable consumer reviews and advertising as complements, and they consider the competitor's favorable reviews as a strategic substitute to their own advertising. Further, firms' optimal price responses to any change in reviews are in the same direction as their respective advertising responses. Mathematically, if $k_i > k_i^{(c)}$ for $i \in \{1, 2\}$, then $\frac{\partial a_i^}{\partial \beta_i} > 0$, $\frac{\partial p_i^*}{\partial \beta_i} > 0$, $\frac{\partial a_i^*}{\partial \beta_j} < 0$, and $\frac{\partial p_i^*}{\partial \beta_j} < 0$.*

A firm's favorable or improved reviews increase consumers' valuation for its product—the same effect as its advertising. One may thus expect the firm with an improvement in its reviews to reduce its advertising especially when advertising is inefficient (i.e., when advertising is very costly or k_i is large). Our analysis shows that this naïve intuition turns out to be wrong. If advertising costs are high, the firm will in fact consider advertising as a complement to its more favorable reviews. That is, the firm will optimally do more advertising when its reviews improve relative to its competitor's. Intuitively, when advertising is costly, the competitor's best response is mainly to reduce its price rather than increasing its advertising to compensate for its unfavorable change in reviews. In fact, the competitor will find it optimal to marginally reduce its advertising level. Such an advertising response by its competitor makes the firm's marginal benefit of advertising higher than its marginal cost; thus the firm will advertise even

more (and raise its price as well). That is, in a competitive market, a firm considers its favorable reviews and advertising to be complements especially when advertising is inefficient and costly.

The flip side of the story is that when advertising is inexpensive (i.e., efficient at raising consumers' willingness to pay), the firm with improved reviews will find it optimal to reduce its advertising to reap the benefit of its reviews. This is because, if advertising costs are very low, the firm's competitor will significantly increase its advertising expenditure to compensate for the change in reviews. As a result, the firm with more favorable reviews will see a lowered marginal benefit of advertising and will optimally reduce advertising rather than increasing it to engage its competitor into an advertising war.

In the new age of much increased consumer control of what to watch and when to watch it, firms' advertising is likely much less effective than in the past. According to the 2009 Nielsen survey of over twenty-five thousand online consumers from fifty countries, 90% of consumers trust product recommendations from people they know and 70% trust consumer opinions posted online. Both percentages are much higher than those for firm-led advertising such as the traditional media of TV, radio, newspaper/magazines, or the Internet media of online ads (video or banner ads, search engine result ads, or text ads on mobile phones).¹⁰ In our model framework, this suggests that a firm will likely find its favorable consumer reviews and advertising to be complements since firm-led advertising tends to be less effective than consumer-generated reviews (implying that advertising is inefficient). Further, if the firm receives favorable consumer-generated reviews, its advertising may actually become more effective as its marketing messages may become more credible in conjunction with favorable external validation by consumer reviews. Intuitively, modeling the firm's advertising efficiency endogenously in this manner will strengthen our result showing the complementarity between a firm's favorable reviews and its advertising. This intuition is indirectly confirmed in Proposition 3 by the examination of comparative statics. When advertising is inefficient, an increase in a firm's advertising efficiency (i.e., a decrease in k_i) will strengthen the complementarities between its favorable reviews and advertising.

¹⁰ <http://blog.nielsen.com/nielsenwire/consumer/global-advertising-consumers-trust-real-friends-and-virtual-strangers-the-most/>, accessed August 2010.

Proposition 3: Effect of k_i on the Impact of Consumer Reviews on Advertising

If advertising is inefficient ($k_i > k_i^{(c)}$), the complementary effect of a firm's favorable reviews on its advertising becomes stronger as its advertising efficiency increases.

We now examine how consumer reviews affect the firm's profit. Generally, a monopolist will always benefit from favorable reviews or any improvement in its reviews since it can take advantage of the corresponding increase in the consumer's willingness to pay. In a competitive market, will a firm necessarily be better off and its competitor worse off if the firm's product reviews improve relative to its competitor's? One may intuitively expect so. However, we show that, interestingly, the opposite may be true.

Proposition 4: Effects of Consumer Reviews on Firm Profits

When advertising is efficient, an improvement in a firm's reviews will hurt its own profit and increase its competitor's profit.

Our analysis shows that when advertising is efficient, a firm's improvement in its consumer reviews will actually hurt its own profitability and benefit its competitor. This sounds very counter-intuitive, so we will closely examine the underlying cause. Suppose that firm i 's product valuation is increased by an amount $\Delta\beta_i V_i$ from its improved reviews. Since advertising is very efficient, firm j will find it effective to simply boost its advertising to compensate for firm i 's improved reviews; its optimal advertising response is an advertising increase of $\Delta a_j^* = \frac{k_i \Delta\beta_i V_i}{k_i + k_j - 18k_i k_j t} > 0$. As a result, firm i 's marginal advertising benefit will drop and its optimal response turns out to be a small reduction in advertising (to avoid an advertising war). Alternatively, one can think of firm i as reaping the benefit of its improved reviews by saving some advertising cost with reduced advertising: $\Delta a_i^* = \frac{-k_j \Delta\beta_i V_i}{k_i + k_j - 18k_i k_j t} < 0$. Without considering firm j 's competitive response, firm i would have been able to increase its price because of its improved reviews. However, because of its competitor's large advertising response, firm i will actually reduce its price at equilibrium: $\Delta p_i^* = \frac{-6k_i k_j t \Delta\beta_i V_i}{k_i + k_j - 18k_i k_j t} < 0$, even though firm j will increase its price:

$\Delta p_j^* = \frac{6k_i k_j t \Delta \beta_i V_i}{k_i + k_j - 18k_i k_j t} > 0$. This is because, as one can easily show, firm j 's price increase is much smaller than the increase in its advertising level: $\Delta p_j^* < \Delta a_j^*$. The net effect on firm i 's market share is straightforwardly computed $\Delta x_{in} = -\frac{3k_i k_j \Delta \beta_i V_i}{k_i + k_j - 18k_i k_j t} < 0$. In summary, if advertising is efficient, an improvement in a firm's reviews can trigger a large advertising increase by its competitor, whose market share and profit will both rise despite its slightly increased price. As a result, the firm with improved reviews actually becomes worse off in terms of profit as it optimally reduces both its advertising and price.

We now compare the equilibrium outcomes of the pre-review and post-review periods to examine how the availability of consumer reviews affects the total advertising expenditure and the price separation in the industry. The total advertising expenditure in the industry is given by $T \equiv \sum g_i(a_i^*) = \sum k_i a_i^{*2}$. When both firms have positive reviews, resulting in consumers' higher willingness to pay for their products, one might expect that there is less need for high advertising expenditure because positive reviews represent a form of external validation of product values. We know that firm decisions are affected by consumer reviews only to the extent that reviews influence firms' product valuations to a different degree. Note that intuitively $\beta_i V_i - \beta_j V_j$ represents the relative quality separation between firms' products. Without loss of generality, we assume that firm i receives more favorable reviews, i.e., $(\beta_i V_i - \beta_j V_j) - (V_i - V_j) > 0$. Therefore, consumer reviews have increased firm i 's perceived product quality relative to firm j 's product. The following propositions shed light on how consumer reviews affect the industry's advertising expenditure and pricing.

Proposition 5: Effect of Consumer Reviews on Total Industry Advertising Expenditure

The larger firm i 's relative improvement in perceived quality due to consumer reviews, i.e., the larger $(\beta_i V_i - \beta_j V_j) - (V_i - V_j)$ is, the more likely the industry advertising expenditure will increase. Further, the more advertising-efficient firm i is (i.e., smaller k_i , or smaller $k_i - k_j$), the more likely the total industry advertising expenditure will rise when reviews become available.

Proposition 6: Effect of Consumer Reviews on Industry Price Difference

Consumer reviews' effect on firms' price separation can be in the opposite direction to their effect on the relative separation of firms' perceived qualities. That is, consumer review information can reduce (increase) the price difference between firms' products even when it increases (reduces) the separation between firms' perceived qualities.

We find that how consumer reviews change the total industry advertising depends on the amount of the favorable firm's improvement in its *relative* quality perception and its relative advertising efficiency. If firm i 's favorable change in its relative quality perception is large and its relative advertising efficiency is high, firm i will more likely find a lower *marginal* advertising benefit. To explain the intuition behind Proposition 5, we consider a specific case of an efficient-advertising market (i.e. advertising is not very costly). Firm i will reduce advertising when it receives more favorable reviews than firm j , because firm j will see an increased marginal advertising benefit (due to its less favorable reviews) and hence increase its advertising. The higher firm i 's improvement in quality perception, $(\beta_i V_i - \beta_j V_j) - (V_i - V_j)$, the lower firm i 's *marginal* advertising benefit and the higher firm j 's *marginal* advertising benefit. The higher firm i 's advertising efficiency (i.e., the lower k_i or $k_i - k_j$), the lower firm i 's *marginal* advertising benefit when it receives favorable reviews, because its advertising level must have already been very high, making it suboptimal for firm i to further increase its advertising. So, the larger improvement in firm i 's relative quality perception and the higher its advertising efficiency, the more incentive (i.e., higher marginal benefit) firm j has to significantly increase its advertising, making it more likely that the total industry advertising expenditure will rise (i.e., firm j increases its advertising by a larger amount than firm i 's reduction in advertising).

One may intuitively expect that, if consumer review information increases the separation in firms' relative perceived qualities (i.e., if $|\beta_i V_i - \beta_j V_j| > |V_i - V_j|$), the price difference between the firm's products will tend to increase as well. Proposition 6 shows that this may not necessarily be the case. Consumer reviews may have opposite effects on firms' quality separation and their price separation. The

main insight here also comes from the competitive advertising responses to consumer reviews. In essence, when consumer reviews increase the separation in the perceived product qualities, the difference in firms' marginal incentives to advertise may help counter that increased separation and can, under many circumstances, yield higher rather than lower competitive pressure on prices.

Next we investigate how changes in consumers' horizontal taste influence the results. How will the strength of the consumers' taste preference affect firms' advertising or the effect of consumer reviews on advertising? The following propositions shed light on these questions.

Proposition 7: Effect of t on the Impact of Consumer Reviews on Advertising

As the strength of consumers' taste preference increases, the effect of consumer reviews on firms' advertising levels becomes stronger (weaker) if advertising is efficient (inefficient).

Proposition 8: Effect of t on Firms' Advertising Levels

As consumers' taste preference becomes stronger (i.e., as t increases), one firm will increase advertising whereas the other will reduce its advertising. The firm with the lower potential variable-margin $\beta_i V_i + e_i^ - c_i - f_i(e_i^*)$ is more likely to increase advertising as t increases.*

Proposition 7 shows that the strength of consumers' taste preference influences the impact of consumer reviews on advertising in a dyadic way depending on the advertising efficiency. As t increases, the effect of consumer reviews on advertising will become stronger when advertising is efficient and weaker when advertising is inefficient. The intuition behind the dyadic nature of these second-order comparative statics comes from the fact that a firm's favorable reviews may be a substitute or complement to its advertising depending on the level of advertising efficiency.

Proposition 8 is more interesting though it is not directly related to consumer reviews. The strength t of consumers' taste preference clearly affects firms' advertising decisions. A larger t has two effects. First, it reduces consumers' willingness to pay because their horizontal preferences are not exactly matched. Second, a larger t implies more differentiation between firms' products hence reducing their price competition, which is manifested in equations (3) and (4). Thus, it is not clear how firms'

advertising levels will change if t becomes larger. One might expect that both firms may have incentives to increase brand advertising to raise consumers' valuations further since a larger t tends to lower both the consumer's product valuations and the price competition. Proposition 8 shows a surprising result—firms' optimal advertising responses to an increase in t are qualitatively different: one firm will increase advertising whereas the other will reduce it. As t becomes larger, the firm with a relatively low variable margin tends to increase advertising whereas its competitor tends to reduce advertising. For example, if both firms are equally effective at advertising ($k_i = k_j$), the firm with a lower variable margin ($\alpha_i V_i + e_i^* - c_i - f_i(e_i^*)$) will have a higher marginal incentive to increase advertising when t becomes larger alleviating price competition. As a result, the firm with a higher variable margin will find it optimal to reduce its advertising rather than to engage in an advertising war.

2.4. Alternative Sequential Decision Making and Implications

Thus far we have assumed that firms make their decisions in the following sequence: advertising, product premium and finally, price. Arguably, firms may make their advertising (fixed-cost) and product premium (variable-cost) decisions together simultaneously rather than sequentially. Thus, a reasonable alternative dynamic game may have the following two stages. In the first stage, each firm simultaneously selects its levels of advertising (a_i) and product premium (e_i). In the second stage, firms choose their prices simultaneously and consumers subsequently make purchases. Again the equilibrium outcome is solved through the backwards induction technique. Since the pricing stage is the same as before, we obtain the same equilibrium price and profit expressions for the pricing subgame.

$$\Pi_i^*(e_1, e_2, a_1, a_2) = \frac{[3t + \beta_i V_i + a_i + e_i - \beta_j V_j - a_j - e_j - c_i + c_j - f_i(e_i) + f_j(e_j)]^2}{18t} - g_i(a_i), i \in \{1, 2\}.$$

In the first stage, each firm chooses both advertising and product premium to maximize its profit. Note that, under the condition (A1) for a competitive equilibrium in the price subgame, firm i 's profit is a monotonically increasing function in $e_i - f_i(e_i)$ for any values of a_i , a_j , or e_j . Thus, firm i 's optimal product premium is $e_i^* = \operatorname{argmax}_{e_i \geq 0} \{e_i - f_i(e_i)\}$, the same equilibrium premium outcome as before.

Therefore, our alternative 2-stage dynamic game must also yield exactly the same advertising equilibrium outcome as before.

We now consider whether our results depend on the sequential nature of the game. Will our results remain robust in a simultaneous (static) game setting, in which both firms set their respective prices, advertising and product premium all simultaneously? Solving for the location of the consumer who is indifferent between the two products (i.e., $U_i = U_j$), we obtain each firm's market share: $m_i = \frac{\beta_i V_i + a_i + e_i - \beta_j V_j - a_j - e_j + t - p_i + p_j}{2t}$. Firm i 's profit function is thus expressed as $\Pi_i(p_i, p_j, e_i, e_j, a_i, a_j) = [p_i - c_i - f_i(e_i)] m_i - k_i a_i^2$. We simultaneously solve six FOCs for profit maximization: $\frac{\partial \Pi_i}{\partial p_i} = 0$, $\frac{\partial \Pi_i}{\partial a_i} = 0$, $\frac{\partial \Pi_i}{\partial e_i} = 0$, $i = 1, 2$. From $\frac{\partial \Pi_i}{\partial p_i} = 0$, $i = 1, 2$, we obtain $p_i = \frac{1}{3}[\beta_i V_i - \beta_j V_j + a_i - a_j + e_i - e_j + 2c_i + 2f_i(e_i) + c_j + f_j(e_j) + 3t]$.

Substituting the p_i and p_j expressions into FOC with respect to e_i , we get

$$\begin{aligned} \frac{\partial \Pi_i}{\partial e_i} &= -f'_i(e_i) \frac{\beta_i V_i + a_i + e_i - \beta_j V_j - a_j - e_j + t - p_i + p_j}{2t} + \frac{[p_i - c_i - f_i(e_i)]}{2t} \\ &= \frac{1}{6t} [3t + \beta_i V_i + a_i + e_i - \beta_j V_j - a_j - e_j - c_i + c_j - f_i(e_i) + f_j(e_j)] [1 - f'_i(e_i)] = 0. \end{aligned}$$

Note that the expression in the first square brackets is positive under the assumption (A1). Thus, this FOC is equivalent to $1 - f'_i(e_i) = 0$, and the optimal level of product premium is the same as $e_i^* = \operatorname{argmax}_{e_i \geq 0} \{e_i - f_i(e_i)\}$. Given e_i^* , we can simultaneously solve the four linear FOCs with respect to p_i and a_i ; the final equilibrium outcome is as follows.

$$\begin{aligned} a_i^* &= \frac{2k_j [3t + \beta_i V_i - \beta_j V_j - c_i + c_j + e_i^* - f_i(e_i^*) - e_j^* + f_j(e_j^*)] - 1}{2(12k_i k_j t - k_i - k_j)} \\ p_i^* &= \frac{2k_i t [2k_j (3t + \beta_i V_i - \beta_j V_j - c_i + c_j + e_i^* - f_i(e_i^*) - e_j^* + f_j(e_j^*)) - 1]}{12k_i k_j t - k_i - k_j} + c_i + f_i(e_i^*) \\ \Pi_i^* &= \frac{k_i (8k_i t - 1) [2k_j (3t + \beta_i V_i - \beta_j V_j - c_i + c_j + e_i^* - f_i(e_i^*) - e_j^* + f_j(e_j^*)) - 1]^2}{4(12k_i k_j t - k_i - k_j)^2}, \end{aligned}$$

We observe that the simultaneous (static) equilibrium outcome has similar functional forms to the sequential-move (dynamic) equilibrium outcome (12)-(14) except that a few numeric coefficients are

different. In fact, all our results from the dynamic game stay qualitatively the same for the static game; only the parameter regions (of inefficient versus efficient advertising) are different. The new advertising efficiency cutoffs are $k_i^{(c)} \equiv \frac{1}{2[3t - \beta_i V_i + \beta_j V_j + c_i - c_j - e_i^* + f_i(e_i^*) + e_j^* - f_j(e_j^*)]}$. The efficient advertising market is defined by $\frac{1}{8t} < k_i < k_i^{(c)}$ for $i \in \{1, 2\}$ while the inefficient advertising market is defined by $k_i > k_i^{(c)}$ for $i \in \{1, 2\}$. Thus, we have shown that our main findings are quite robust to the extensive form of the game.

2.5. Strategic Manipulation of Reviews

Anecdotal evidence suggests that sellers spend resources to manipulate their product reviews (ranging from directly creating reviews to paying their early customers to write favorable reviews). With some circumstantial evidence of manipulation, Kornish (2009) finds that manipulation of reviews is widespread though not dominant. We now extend our base model to study firms' incentives to manipulate their product reviews. Firms may spend resources on their early customers (e.g., offering better services) to foster better reviews or they may directly manipulate their product reviews. We modify the game we studied earlier as follows. In the first stage of the game, firms make simultaneous decisions on their own levels of advertising (a_i) and of strategic manipulation of their reviews (s_i). The second and third stages are the same as before.

Firm i 's manipulation ($s_i \geq 0$) of its product reviews will serve to increase consumers' *ex ante* expected utility for its product; we define s_i to be the amount of *ex ante* valuation increase. Thus, given firm i 's advertising, manipulation of reviews, product premium and price, consumers of type x will derive a surplus of $U_i = \beta_i V_i + a_i + e_i + s_i - tx_i - p_i$ from product i , where x_i is the distance between firm i and consumer x .¹¹ We assume that both firms are equally effective at manipulating their reviews; similar

¹¹ We can alternatively define s_i such that $U_i = (\beta_i + s_i)V_i + a_i + e_i - x_i t - p_i$. Our current definition differs from this by a scalar factor; it allows us to demonstrate the parallel between the manipulation of reviews and advertising. Further, one can also define s_i to be a combination of the increase in firm i 's own valuation and the reduction of the competitor's product valuation (e.g., by bad mouthing the competitor's product). These alternative definitions will not change our results.

to the cost of advertising, the cost of manipulating reviews is also assumed to be quadratic in the level of manipulation: bs_i^2 . Firm i 's profit is given by

$\Pi_i(p_i, p_j, e_i, e_j, a_i, a_j, s_i, s_j) = [p_i - c_i - f_i(e_i)]m_i - k_i a_i^2 - bs_i^2$, where firm i 's market share m_i is given by $m_i = \frac{\beta_i V_i + a_i + e_i + s_i - \beta_j V_j - a_j - e_j - s_j + t - p_i + p_j}{2t}$. Backwards induction for the last two stages is the same as before except for the additional factor s_i ; we obtain the profit functions for the first stage of the game:

$$\Pi_i^*(a_i, a_j, s_i, s_j) = \frac{[3t + \beta_i V_i + a_i + s_i + e_i^* - \beta_j V_j - a_j - s_j - e_j^* - c_i + c_j - f_i(e_i^*) + f_j(e_j^*)]^2}{18t} - k_i a_i^2 - bs_i^2.$$

As before, we are interested in the interior solutions for competitive equilibrium. The four FOCs, $\frac{\partial \Pi_i^*}{\partial a_i} = 0$

and $\frac{\partial \Pi_i^*}{\partial s_i} = 0$ for $i = 1, 2$, are simultaneously solved to yield the following:

$$a_i^* = \frac{k_j + b \{1 - 3k_j [3t + \beta_i V_i - \beta_j V_j - c_i + c_j + e_i^* - f_i(e_i^*) - e_j^* + f_j(e_j^*)]\}}{3[2k_i k_j - b(18k_i k_j t - k_i - k_j)]},$$

$$s_i^* = \frac{k_i k_j + k_i b \{1 - 3k_j [3t + \beta_i V_i - \beta_j V_j - c_i + c_j + e_i^* - f_i(e_i^*) - e_j^* + f_j(e_j^*)]\}}{3b[2k_i k_j - b(18k_i k_j t - k_i - k_j)]}$$

Proposition 9: Manipulation of Consumer Reviews

- (a) *A firm's best advertising response and its best strategic manipulation response to a change in its reviews are in the same direction.*
- (b) *The firm with a favorable exogenous change in its reviews is more likely to increase (reduce) its manipulation whereas its competitor is more likely to reduce (increase) its manipulation when the cost of manipulation is high (low).*

The first part of Proposition 9 states that when a firm's reviews change, it will increase or decrease both its advertising and its manipulation of reviews. This is reasonable since both advertising and review manipulation alter consumer valuations by incurring a fixed cost. Proposition 9 also shows that a firm is more likely to complement its favorable reviews with additional manipulation of its reviews when the cost of manipulation is *high* than when it is *low*. The intuition here is similar to that of

advertising in Proposition 2. We can essentially interpret advertising as the firm's attempt to influence or manipulate its consumer reviews. Our results suggest that the firm receiving more favorable reviews will in fact have more incentives to manipulate its reviews when the manipulation of reviews is costly (or inefficient). When the cost of manipulation is low, the firm with more favorable reviews will reduce its manipulation because of the anticipated response of its competitor's increased manipulation.

2.6. Limitations

As with any analytical model, our model has some limitations. First, we have assumed that positive reviews, advertising, and product premium are perfect substitutes in an additive manner in terms of influencing the consumer's product valuation. It is possible that there might be non-linear interactions between these factors. In particular, how effective a firm's advertising is in terms of increasing consumers' valuation may depend on how favorable its reviews are; one might expect a positive correlation between the two. In other words, consumer reviews may endogenously influence firms' advertising effectiveness. For example, if a firm receives favorable reviews, its advertising may become more effective or its competitor's advertising less effective. Accommodating such non-linear interactions in our model makes it analytically intractable. However, intuitively, it tends to *strengthen* our complementarity result between favorable reviews and advertising because the firm with favorable reviews will find a larger marginal return of advertising whereas its competitor a smaller marginal return. We expect our main results to remain qualitatively the same in such a model extension.

Second, we have assumed that reviews influence the consumers uniformly—changing all consumers' willingness to pay by the same amount. It is not unlikely that consumers may have different valuations for quality, perceive the reviews somewhat differently, or have different levels of knowledge about the product. That is, it is possible that reviews may change some consumers' valuations more than others. However, we expect that our main results will remain qualitatively the same even if we introduce some heterogeneity in consumers' responses to reviews, e.g., if there are two types of consumers, one of which (perhaps an experienced or expert consumer) is not affected by reviews while the other is. We will

rely on future research to find any additional insights that other types of consumer heterogeneity may bring forth.

Third, we do not model any potential effect of “advertising hype” on consumer reviews. In particular, it is possible that too much advertising in the early (pre-review) period might build very high expectations leading to more negative reviews in the later period from disappointed customers. However, such possibilities will actually *strengthen* our main complementarity result. With such negative interactions, it is even *more* likely that the firm with favorable reviews will *reduce* its advertising in the current period, such that in the next period, its product reviews (from additional consumers) will not suffer too much from the advertising hype in the previous period. Lastly, in our model, each firm offers only one product with a single level of product premium. Allowing for multiple product offerings may change firms’ strategic incentives to respond to consumer reviews. We hope that our research motivates further research on firms’ strategic marketing responses to UGC.

2.7. Conclusion

Consumer product reviews, one of the earliest forms of online UGC, have now become very prevalent in online retailing and play a very important role in consumers’ purchase decisions. These reviews help consumers resolve or reduce uncertainties about product features or qualities prior to their purchases and thereby allow for better and more informed decisions. We have provided a normative, game-theoretic model to study how consumer reviews influence competitive firms’ advertising, product premium, and pricing strategies. Our research shows several surprising and interesting findings. First, consumer reviews and firms’ advertising do not affect firms’ product premium strategies. Second, even though a firm’s favorable consumer reviews and advertising are perfect substitutes with respect to increasing the consumer’s willingness to pay, the firm may actually consider them as complements, especially when advertising is expensive. The complementarities between the firm’s favorable reviews and its advertising will strengthen if its favorable reviews also make its advertising more effective. Our analysis shows that competitive firms have qualitatively different advertising responses to the availability

of reviews, i.e., one firm will increase advertising whereas the other will decrease its advertising. Third, the total industry advertising expenditure may increase even when both firms have positive reviews that increase consumers' product valuations. Fourth, counterintuitively, the effect of consumer reviews on the difference between firms' prices can be in the *opposite* direction to their effect on the relative separation between firms' perceived qualities. Fifth, surprisingly, an improvement in a firm's reviews may hurt its own profit and increase its competitor's profit. Sixth, the strength of consumers' taste preference influences the impact of consumer reviews on advertising and firms have qualitatively different best advertising responses to a stronger consumer taste preference. Lastly, a firm is more likely to complement its favorable reviews with additional manipulation of its reviews when the cost of manipulation is high than when it is low.

Appendix for Chapter 2

Proof of Proposition 1:

Firms will set their respective e_i to maximize their profits:

$$\Pi_i^*(e_1, e_2, a_1, a_2) = \frac{[3t + \beta_i V_i + a_i + e_i - \beta_j V_j - a_j - e_j - c_i + c_j - f_i(e_i) + f_j(e_j)]^2}{18t} - g_i(a_i), \quad i \in \{1, 2\}.$$

The necessary condition (A1) for a competitive pricing equilibrium requires that

$$3t + \beta_i V_i + a_i + e_i - \beta_j V_j - a_j - e_j - c_i + c_j - f_i(e_i) + f_j(e_j) > 0.$$

Thus, $\Pi_i^*(e_1, e_2, a_1, a_2)$ is a monotonically increasing function in $e_i - f_i(e_i)$. Therefore, to maximize

$\Pi_i^*(e_1, e_2, a_1, a_2)$, we maximize $e_i - f_i(e_i)$; that is, $e_i^* = \operatorname{argmax}_{e_i \geq 0} \{e_i - f_i(e_i)\}$. \square

Proof of positive advertising levels and profits at advertising equilibrium:

We first show the following lemma.

Lemma 1: $18k_i k_j t - k_i - k_j > 0$ if $k_i > k_i^{(c)}$ and $k_j > k_j^{(c)}$;

$$18k_i k_j t - k_i - k_j < 0 \text{ if } \frac{1}{18t} < k_i < k_i^{(c)} \text{ and } \frac{1}{18t} < k_j < k_j^{(c)}.$$

Proof: Define a function: $Z(k_i, k_j) \equiv 18k_i k_j t - k_i - k_j$. Z is an increasing function in both k_i and k_j

since $\frac{\partial Z}{\partial k_i} = 18k_j t - 1 > 0$ and $\frac{\partial Z}{\partial k_j} = 18k_i t - 1 > 0$. Note that at $k_i = k_i^{(c)}$ and $k_j = k_j^{(c)}$, $Z(k_i, k_j) =$

$$0. \text{ Thus, } Z(k_i, k_j) = \begin{cases} > 0, & \text{if } k_i > k_i^{(c)} \text{ and } k_j > k_j^{(c)}; \\ < 0, & \text{if } \frac{1}{18t} < k_i < k_i^{(c)} \text{ and } \frac{1}{18t} < k_j < k_j^{(c)}. \end{cases} \quad \text{QED.}$$

With Lemma 1, it is obvious that in both the efficient and the inefficient advertising markets, $a_i^* > 0$ and

$\Pi_i^* > 0$ for both firms. □

Proof of Proposition 2:

Using Lemma 1, it is straightforward to see that when $k_i > k_i^{(c)}$ and $k_j > k_j^{(c)}$,

$$\frac{\partial a_i^*}{\partial \beta_i} = \frac{k_j V_i}{18k_i k_j t - k_i - k_j} > 0, \quad \frac{\partial p_i^*}{\partial \beta_i} = \frac{6k_i k_j V_i t}{18k_i k_j t - k_i - k_j} > 0,$$

$$\frac{\partial a_i^*}{\partial \beta_j} = \frac{-k_j V_j}{18k_i k_j t - k_i - k_j} < 0, \quad \frac{\partial p_i^*}{\partial \beta_j} = \frac{-6k_i k_j V_j t}{18k_i k_j t - k_i - k_j} < 0. \quad \square$$

Proof of Proposition 3:

In Proposition 2, we have shown the complementarities between a firm's favorable reviews and its

advertising when $k_i > k_i^{(c)}$ for $i = 1, 2$: $\frac{\partial a_i^*}{\partial \beta_i} = \frac{k_j V_i}{18k_i k_j t - k_i - k_j} > 0$. Now we examine the second-order

comparative statics:

$$\frac{\partial}{\partial k_i} \left(\frac{\partial a_i^*}{\partial \beta_i} \right) = \frac{\partial}{\partial k_i} \left(\frac{k_j V_i}{18k_i k_j t - k_i - k_j} \right) = \frac{-k_j V_i (18k_j t - 1)}{(18k_i k_j t - k_i - k_j)^2} < 0.$$

Thus, if a firm's favorable reviews also make its advertising more efficient (i.e., k_i becomes smaller), the

complementarities between the firm's favorable reviews and its advertising will strengthen. □

Proof of Proposition 4:

$$\frac{\partial \Pi_i^*}{\partial \beta_i} = \frac{2k_i k_j V_i (18k_i t - 1) \{ 3k_j [3t + \beta_i V_i - \beta_j V_j - c_i + c_j + e_i^* - f_i(e_i^*) - e_j^* + f_j(e_j^*)] - 1 \}}{3(18k_i k_j t - k_i - k_j)^2}$$

$$\frac{\partial \Pi_j^*}{\partial \beta_i} = \frac{2k_i k_j V_j (18k_j t - 1) \{ 3k_i [3t - \beta_i V_i + \beta_j V_j + c_i - c_j - e_i^* + f_i(e_i^*) + e_j^* - f_j(e_j^*)] - 1 \}}{3(18k_i k_j t - k_i - k_j)^2}$$

When advertising is efficient (i.e., $\frac{1}{18t} < k_i < k_i^{(c)}$), we obtain $\frac{\partial \Pi_i^*}{\partial \beta_i} < 0$ and $\frac{\partial \Pi_j^*}{\partial \beta_i} > 0$. \square

Proof of Proposition 5:

We use i to denote the firm that has more favorable reviews, i.e., $(\beta_i V_i - \beta_j V_j) - (V_i - V_j) > 0$. With equation (12), the firm's advertising expenditure is given by

$$g_i(a_i^*) = k_i a_i^{*2} = \frac{k_i \{3k_j(3t + \alpha_i V_i - \alpha_j V_j - c_i + c_j + e_i^* - f_i(e_i^*) - e_j^* + f_j(e_j^*)) - 1\}^2}{9(18k_i k_j t - k_i - k_j)^2}.$$

Let Δg_i denote firm i 's change in advertising expenditure when consumer reviews are available relative to when reviews are absent.

$$\begin{aligned} \Delta g_i &= \frac{k_i \{3k_j(3t + \beta_i V_i - \beta_j V_j - c_i + c_j + e_i^* - f_i(e_i^*) - e_j^* + f_j(e_j^*)) - 1\}^2}{9(18k_i k_j t - k_i - k_j)^2} - \frac{k_i \{3k_j(3t + V_i - V_j - c_i + c_j + e_i^* - f_i(e_i^*) - e_j^* + f_j(e_j^*)) - 1\}^2}{9(18k_i k_j t - k_i - k_j)^2} \\ &= \frac{k_i k_j [\beta_i V_i - \beta_j V_j - (V_i - V_j)] \{6k_j(3t - c_i + c_j + e_i^* - f_i(e_i^*) - e_j^* + f_j(e_j^*)) + 3k_j(\beta_i V_i - \beta_j V_j + V_i - V_j) - 2\}}{3(18k_i k_j t - k_i - k_j)^2} \end{aligned}$$

The industry's change in advertising expenditure is given by $\Delta T = \Delta g_i + \Delta g_j$

$$\begin{aligned} &= \frac{k_i k_j [\beta_i V_i - \beta_j V_j - (V_i - V_j)] \{6k_j(3t - c_i + c_j + e_i^* - f_i(e_i^*) - e_j^* + f_j(e_j^*)) + 3k_j(\beta_i V_i - \beta_j V_j + V_i - V_j) - 2\}}{3(18k_i k_j t - k_i - k_j)^2} \\ &+ \frac{k_i k_j [\beta_j V_j - \beta_i V_i - (V_j - V_i)] \{6k_i(3t - c_j + c_i + e_j^* - f_j(e_j^*) - e_i^* + f_i(e_i^*)) + 3k_i(\beta_j V_j - \beta_i V_i + V_j - V_i) - 2\}}{3(18k_i k_j t - k_i - k_j)^2} \\ &= \frac{k_i k_j [\beta_i V_i - \beta_j V_j - (V_i - V_j)] \{6t(k_j - k_i) + 2(k_i + k_j)(-c_i + c_j + e_i^* - f_i(e_i^*) - e_j^* + f_j(e_j^*)) + (k_i + k_j)(\beta_i V_i - \beta_j V_j + V_i - V_j)\}}{(18k_i k_j t - k_i - k_j)^2} \end{aligned}$$

Thus, $\Delta T > 0$ if $\{6t(k_j - k_i) + 2(k_i + k_j)(-c_i + c_j + e_i^* - f_i(e_i^*) - e_j^* + f_j(e_j^*)) + (k_i + k_j)(\beta_i V_i - \beta_j V_j + V_i - V_j)\} > 0$. The latter inequality condition is equivalent to

$$\beta_i V_i - \beta_j V_j - (V_i - V_j) > \frac{6t(k_i - k_j)}{k_i + k_j} + 2[V_j - V_i + c_i - c_j - e_i^* + f_i(e_i^*) + e_j^* - f_j(e_j^*)].$$

This condition is more likely to hold for larger $\beta_i V_i - \beta_j V_j - (V_i - V_j)$ and smaller k_i . That is, the larger firm i 's relative improvement in perceived quality due to consumer reviews and the more advertising-

efficient firm i is, the more likely the industry advertising expenditure will increase as a result of consumer reviews. □

Proof of Proposition 6:

Firm i is assumed to have more favorable reviews: $(\beta_i V_i - \beta_j V_j) - (V_i - V_j) > 0$.

We also assume that consumer reviews will not create any crossover in price or perceived qualities. That is, if a firm has a higher baseline perceived quality or price than its competitor when there are no reviews, it will also have higher a baseline perceived quality or price when reviews are available. Intuitively, this implies that the effects of consumer reviews on product valuations are not drastic, e.g., a higher priced, higher (*ex ante* perceived) quality BMW 5-series car will still have a higher price and higher perceived quality than a Honda Accord car even if the Honda Accord receives more favorable reviews. We look at two cases.

Case 1: $V_i > V_j$

With the presence of consumer review information, the perceived quality separation between firms' products has *increased* since $\beta_i V_i - \beta_j V_j > V_i - V_j > 0$. If $p_i^{*(o)} > p_j^{*(o)}$, then the price difference between the products will *decrease* in an *efficient*-advertising market (with $\frac{1}{18t} < k_i < k_i^{(c)}$ and $\frac{1}{18t} < k_j < k_j^{(c)}$). This is because p_i^* will decrease from $p_i^{*(o)}$ whereas p_j^* will increase from $p_j^{*(o)}$ since $\frac{\partial p_i^*}{\partial(\beta_i V_i - \beta_j V_j)} = \frac{1}{V_i} \frac{\partial p_i^*}{\partial \beta_i} < 0$ and $\frac{\partial p_j^*}{\partial(\beta_i V_i - \beta_j V_j)} = \frac{1}{V_i} \frac{\partial p_j^*}{\partial \beta_i} > 0$ when advertising is efficient. Similarly, if $p_i^{*(o)} < p_j^{*(o)}$, then the price difference between the products will *decrease* in an *inefficient*-advertising market (with $k_i > k_i^{(c)}$ and $k_j > k_j^{(c)}$). This is because p_i^* will increase from $p_i^{*(o)}$ whereas p_j^* will decrease from $p_j^{*(o)}$ since $\frac{\partial p_i^*}{\partial(\beta_i V_i - \beta_j V_j)} > 0$ and $\frac{\partial p_j^*}{\partial(\beta_i V_i - \beta_j V_j)} < 0$ when advertising is inefficient.

Case 2: $V_i < V_j$

In this case, with the presence of consumer review information, the perceived quality separation between firms' products has *decreased* since $\beta_j V_j - \beta_i V_i < V_j - V_i > 0$. If $p_i^{*(o)} > p_j^{*(o)}$, then the price

difference between the products will *increase* in an *inefficient*-advertising market (with $k_i > k_i^{(c)}$ and $k_j > k_j^{(c)}$). This is because p_i^* will increase from $p_i^{*(o)}$ whereas p_j^* will decrease from $p_j^{*(o)}$ since $\frac{\partial p_i^*}{\partial(\beta_i V_i - \beta_j V_j)} = \frac{1}{V_i} \frac{\partial p_i^*}{\partial \beta_i} > 0$ and $\frac{\partial p_j^*}{\partial(\beta_i V_i - \beta_j V_j)} = \frac{1}{V_i} \frac{\partial p_j^*}{\partial \beta_i} < 0$ when advertising is inefficient. Similarly, if $p_i^{*(o)} < p_j^{*(o)}$, then the price difference between the products will *increase* in an *efficient*-advertising market (with $\frac{1}{18t} < k_i < k_i^{(c)}$ and $\frac{1}{18t} < k_j < k_j^{(c)}$). This is because p_i^* will decrease from $p_i^{*(o)}$ whereas p_j^* will increase from $p_j^{*(o)}$ since $\frac{\partial p_i^*}{\partial(\beta_i V_i - \beta_j V_j)} < 0$ and $\frac{\partial p_j^*}{\partial(\beta_i V_i - \beta_j V_j)} > 0$ when advertising is efficient.

In conclusion, we have shown that under the circumstances discussed above, consumer review information can *reduce* (increase) the price difference between firms' products even though it *increases* (reduces) the separation between firms' perceived qualities. \square

Proof of Proposition 7:

$$\frac{\partial}{\partial t} \left(\frac{\partial a_i^*}{\partial \beta_i} \right) = \frac{\partial}{\partial t} \left(\frac{k_j V_i}{18k_i k_j t - k_i - k_j} \right) = \frac{-18k_i k_j^2 V_i}{(18k_i k_j t - k_i - k_j)^2} < 0.$$

$$\frac{\partial}{\partial t} \left(\frac{\partial a_i^*}{\partial \beta_j} \right) = \frac{\partial}{\partial k_j} \left(\frac{-k_j V_j}{18k_i k_j t - k_i - k_j} \right) = \frac{18k_i k_j^2 V_j}{(18k_i k_j t - k_i - k_j)^2} > 0.$$

From Table 2.1, we know

$$\frac{\partial a_i^*}{\partial \beta_i} = \begin{cases} > 0, & \text{if } k_i > k_i^{(c)} \text{ and } k_j > k_j^{(c)}; \\ < 0, & \text{if } \frac{1}{18t} < k_i < k_i^{(c)} \text{ and } \frac{1}{18t} < k_j < k_j^{(c)}; \end{cases}$$

$$\frac{\partial a_i^*}{\partial \beta_j} = \begin{cases} > 0, & \text{if } \frac{1}{18t} < k_i < k_i^{(c)} \text{ and } \frac{1}{18t} < k_j < k_j^{(c)}; \\ < 0, & \text{if } k_i > k_i^{(c)} \text{ and } k_j > k_j^{(c)}; \end{cases}$$

Thus, if $\frac{1}{18t} < k_i < k_i^{(c)}$ and $\frac{1}{18t} < k_j < k_j^{(c)}$, then $\frac{\partial}{\partial t} \left(\left| \frac{\partial a_i^*}{\partial \beta_i} \right| \right) > 0$ and $\frac{\partial}{\partial t} \left(\left| \frac{\partial a_i^*}{\partial \beta_j} \right| \right) > 0$; if $k_i > k_i^{(c)}$ and

$k_j > k_j^{(c)}$, then $\frac{\partial}{\partial t} \left(\left| \frac{\partial a_i^*}{\partial \beta_i} \right| \right) < 0$ and $\frac{\partial}{\partial t} \left(\left| \frac{\partial a_i^*}{\partial \beta_j} \right| \right) < 0$. That is, as t increases, the effect of consumer reviews

on firms' advertising levels becomes stronger (weaker) if advertising is efficient (inefficient). \square

Proof of Proposition 8:

$$\frac{\partial a_i^*}{\partial t} = \frac{3k_j \{k_i - k_j - 6k_i k_j [\beta_i V_i - \beta_j V_j - c_i + c_j + e_i^* - f_i(e_i^*) - e_j^* + f_j(e_j^*)]\}}{(18k_i k_j t - k_i - k_j)^2}.$$

Switching the labels i and j , we get

$$\frac{\partial a_j^*}{\partial t} = - \frac{3k_i \{k_i - k_j - 6k_i k_j [\beta_i V_i - \beta_j V_j - c_i + c_j + e_i^* - f_i(e_i^*) - e_j^* + f_j(e_j^*)]\}}{(18k_i k_j t - k_i - k_j)^2}.$$

Clearly, the effect of t on advertising is in opposite directions for the two firms.

Thus, $\frac{\partial a_i^*}{\partial t} > 0$ and $\frac{\partial a_j^*}{\partial t} < 0$ if $k_i - k_j - 6k_i k_j [\beta_i V_i - \beta_j V_j - c_i + c_j + e_i^* - f_i(e_i^*) - e_j^* + f_j(e_j^*)] > 0$;

$\frac{\partial a_i^*}{\partial t} < 0$ and $\frac{\partial a_j^*}{\partial t} > 0$ if $k_i - k_j - 6k_i k_j [\beta_i V_i - \beta_j V_j - c_i + c_j + e_i^* - f_i(e_i^*) - e_j^* + f_j(e_j^*)] < 0$.

We observe that the firm with the lower potential variable-margin $\beta_i V_i + e_i^* - c_i - f_i(e_i^*)$ is more likely to increase advertising as t increases. \square

Proof of Proposition 9:

$$\frac{\partial a_i^*}{\partial \beta_i} = \frac{-b k_j V_i}{2k_i k_j - b(18k_i k_j t - k_i - k_j)}$$

$$\frac{\partial s_i^*}{\partial \beta_i} = \frac{-k_i k_j V_i}{2k_i k_j - b(18k_i k_j t - k_i - k_j)}$$

Clearly, if $\frac{1}{18t} < k_i < k_i^{(c)}$ and $\frac{1}{18t} < k_j < k_j^{(c)}$ (i.e., advertising is efficient), then $\frac{\partial a_i^*}{\partial \beta_i} < 0$ and $\frac{\partial s_i^*}{\partial \beta_i} < 0$.

When $k_i > k_i^{(c)}$ and $k_j > k_j^{(c)}$ (i.e., advertising is inefficient), then $\frac{\partial a_i^*}{\partial \beta_i} < 0$ and $\frac{\partial s_i^*}{\partial \beta_i} < 0$ if $b < \frac{2k_i k_j}{18k_i k_j t - k_i - k_j}$, and $\frac{\partial a_i^*}{\partial \beta_i} > 0$ and $\frac{\partial s_i^*}{\partial \beta_i} > 0$ if $b > \frac{2k_i k_j}{18k_i k_j t - k_i - k_j}$. We observe that $\frac{\partial a_i^*}{\partial \beta_i}$ and $\frac{\partial s_i^*}{\partial \beta_i}$ have the same

sign, and that $\frac{\partial s_i^*}{\partial \beta_i}$ is more likely to be positive when b is large than when b is small and vice versa. \square

3. Firm Strategies in the “Mid Tail” of Platform-Based Retailing

3.1. Introduction

Amazon, as a dominant platform-based retailer, not only sells products directly, but also allows hundreds of thousands of third-party sellers (also known as independent sellers) to sell products on its retail platform. Consequently, it offers spectacular range and variety; e.g., it lists for sale over two million products in the “Electronics” category alone. The product variety available on Amazon.com dwarfs what is available at Walmart, the largest traditional (non-platform) retailer, by several orders of magnitude. For example, during April 2010, a staggering 8,010 digital camera products were listed for sale on Amazon whereas 408 such products were offered on Walmart.com and only 30 in a typical, physical Walmart store. Leaving aside the bestsellers, most products available online have low sales, but together they account for a significant portion of Amazon’s total revenue. This phenomenon, popularly known as the “long tail” of internet sales, has been widely documented (Anderson 2006, Brynjolfsson et al. 2003, 2006).

Interestingly, Amazon itself sells only a small percentage of all products listed on its website; most products are sold by third-party sellers. For instance, Amazon directly sells only 7% of the products in its “Electronics” category with the remaining 93% sold by independent sellers. Table 3.1 (second column) shows a similar sales pattern for various other product categories. Third-party sellers can list their products on Amazon.com, which displays these listings to a consumer whenever she conducts a related search.¹² For every unit sold, Amazon charges the seller a fee. In this manner, the third-party sellers benefit from access to the tens of millions of consumers on Amazon.com. In turn, Amazon benefits from these sellers’ sales and the increased product variety, which helps Amazon attract and retain more online customers. Due to these symbiotic advantages, an increasing number of large retailers are establishing similar online retail platforms. For example, Sears has recently launched “Marketplace at Sears.com” to facilitate sales by independent sellers. Clearly, third-party selling on online retail platforms has become an important phenomenon, especially for long-tail products.

¹² For expositional ease, we will refer to the seller as “he,” the consumer as “she,” and the platform owner (Amazon) as “it” throughout the chapter.

Table 3.1: Percentage of products sold by Amazon in two sample product categories

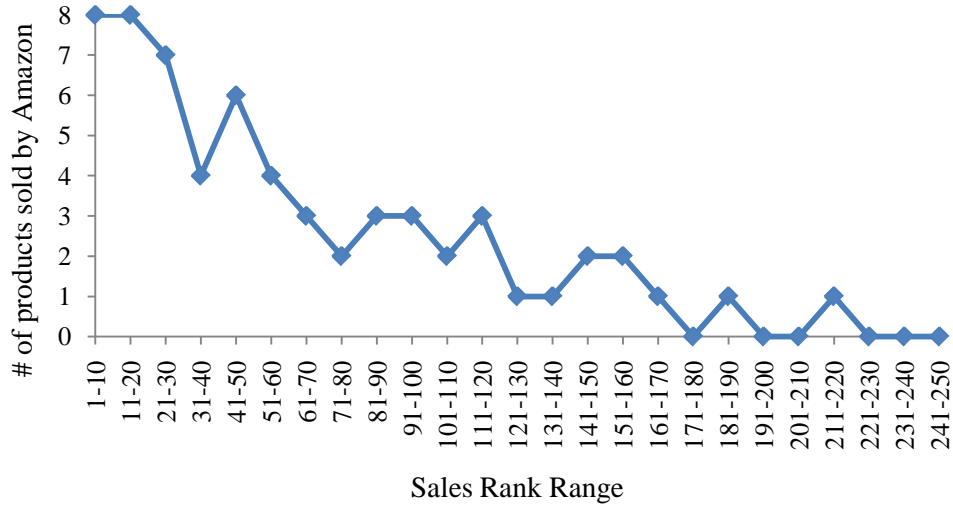
Category/sub-category	Total # of products	% sold by Amazon	% sold by Amazon among top 100 bestsellers
Electronics	2,024,750	7.0	64
--Accessories & Supplies	407,149	10.5	62
--Camera & Photo	410,312	10.1	76
--Car Electronics	16,731	23.3	90
--Computers & Accessories	997,543	4.9	73
--GPS & Navigation	8,453	21.9	89
--Home Audio & Theater	10,433	24.2	71
--Marine Electronics	593	41.1	83
--Office Electronics	39,214	6.7	77
--Portable Audio & Video	48,678	15.1	47
--Security & Surveillance	11,320	15.9	66
--Televisions & Video	14,753	6.4	75
Tools & Home Improvement	2,460,108	5.8	88
--Appliances	12,911	3.7	30
--Building Supplies & Heavy Equipment	147,335	2.0	70
--Fire Safety & Home Security	99,178	6.8	80
--Hardware	434,976	5.5	79
--Heating & Cooling	28,303	6.2	71
--Lighting & Electrical	529,509	3.9	79
--Outdoor Power & Lawn Equipment	61,593	18.3	91
--Painting Tools & Supplies	46,297	6.4	87
--Plumbing	468,885	2.5	72
--Pools, Spas & Supplies	20,770	7.9	62
--Power & Hand Tools	356,371	17.6	92
--Storage & Home Organization	21,940	12.6	88
Sports & Outdoors	3,695,634	3.1	76
Jewelry	1,287,098	3.2	34
Toys & Games	798,977	5.9	66
Shoes	344,710	16.7	72

Source: Data collected on Amazon.com during April 2010

With tens of millions of products available on Amazon, which ones should it procure and sell directly and which ones should it leave to independent sellers to sell? By allowing an independent seller to sell a product, Amazon captures only a fraction of the potential profit. However, given the fixed costs involved in selling a product, selling low-volume items may not be profitable for Amazon. On the other hand, for *specific* niche products, an entrepreneurial and enterprising independent seller might face lower fixed costs and may already have more information than Amazon. (Amazon has data from the sales of millions of products, which it can use to identify high-potential products. However, there may still be niche products for which an independent seller may have better information on demand as compared to Amazon. Anecdotal evidence that we provide subsequently shows that this is a significant phenomenon.) In this context, Amazon's proclivity is to directly sell high-volume products and leave the low-volume items to independent sellers. (The strategy is analogous to that of chain stores wherein the firm itself operates the lucrative city stores but allows franchisees to operate the less attractive dispersed suburban and exurban outlets.) Amazon's strategy on high-volume bestsellers and low-volume long-tail products is rather obvious—it will directly sell the high-volume products and rely on the independent sellers for long-tail products. However, for “mid-tail” products, those that it cannot classify with certainty as either high-volume products or low-volume products, Amazon's strategy is less clear. While Amazon may let independent sellers offer such mid-tail products, it may also be tempted to offer them directly, especially if they show the promise to become bestsellers.

A closer examination of product sales on Amazon's platform confirms the above intuition—Amazon indeed sells a disproportionately large number of high-demand products. For example, though Amazon directly sells only 7% of all electronics products, it sells 64 of the top 100 bestsellers. Table 3.1 (third column) shows that this is consistently true for other product categories. Further, the percentage of products sold directly by Amazon decreases sharply as we go down the list of bestsellers. Figure 3.1 shows an example of this for the “Digital SLR” camera subcategory. In April 2010, this category had 928 products listed; Amazon carried 16 of the top 20 bestsellers, but only 5 of the products with sales ranks from 150 to 250.

Figure 3.1: Bestsellers sold by Amazon in the “Digital SLRs” category



These statistics further suggest that Amazon seems to “cherry-pick” relatively high-demand products from a significant range of mid-tail products for which the *ex ante* expected demand is not sufficiently high for Amazon to readily sell directly, but also not sufficiently low to ignore completely.¹³ Interesting strategic interactions between Amazon and the independent sellers emerge from the uncertainty about the potential demand of these mid-tail products. For a mid-tail product whose sales potential is not readily obvious but which can be sold over a significantly long time horizon, Amazon can initially let the independent seller sell it, track the early sales of the product, and then decide whether or not to offer the product directly. And therein lies the inherent risk faced by a mid-tail independent seller: If the product sells well, Amazon can observe this (since it processes all sales orders on its website) and will likely procure and sell the product directly. When Amazon starts selling the product directly, it can boost its own sales in various ways. For instance, it can prominently display its own offering, and given its advantages in scale and not having to pay its own sales fee, it typically offers lower prices with very competitive or free shipping. Anecdotal evidence from popular online blogs and news sources indicates that Amazon indeed cherry-picks the high-volume products “*in store after store and category after*

¹³ We build the intuition here by considering sales rank (which is based on sales volume) rather than profit. We do this because Amazon publicly releases sales ranks but not any product-level profit data. In our formal model the platform owner makes its decisions based on profit, which is affected by various factors such as demand levels, the marginal cost of procurement and the price sensitivity for the product.

*category, where top-selling products once sold by others are now taken over by Amazon.”*¹⁴ Once Amazon directly procures and sells a product, it will essentially “*take all of the sales away from the [independent] seller.*”¹⁵

This creates a dilemma for the high-demand seller. He may make more profits early on by selling a high volume of the product, but then if Amazon learns that this product is worth selling directly, the seller will lose substantial future sales. Thus, if the seller has a high-demand product, he may have an incentive to reduce his sales to avoid Amazon’s cherry-picking of his products. Anecdotal evidence suggests that some sellers strategically reduce their sales by lowering their services or inventory levels.¹⁶ For instance, they may devote less time and resources to dealing with consumers’ inquiries about their high-demand product or related post-sale services (e.g., they may offer less customization services such as gift wrapping, or answer product inquiries less conscientiously and with a longer time lag). These sellers may also carry a lower inventory level and periodically create stock-out situations.¹⁷ Such service interactions with the consumer typically occur outside Amazon’s retail platform and cannot be directly observed by Amazon. Moreover, with hundreds of thousands of independent sellers, Amazon may find it too costly to monitor even the somewhat observable aspects of seller services. Hence, Amazon may face a demand-learning problem for mid-tail products—if it observes not-so-high unit sales for the seller’s product, it may not be able to infer whether or not the product has the potential for high-enough sales to

¹⁴ <http://www.fool.com/investing/general/2010/02/09/amazons-world-domination-plan.aspx> (accessed 6/1/10)

¹⁵ <http://www.adam-mcfarland.net/2009/10/18/amazon-lets-us-pay-them-to-grow> (accessed 6/1/10)

¹⁶ <http://www.adam-mcfarland.net/2010/07/06/how-amazon-exploits-the-mid-tail> (accessed 10/01/10)

¹⁷ We conducted a small empirical exercise to check for evidence of sellers manipulating their service levels. We selected eight third-party sellers on Amazon selling “mid tail” products (with sales ranks in the upper middle range). We picked different digital cameras and rice cookers, none of which were directly sold by Amazon at that time. Using two different customer email accounts, we emailed each seller two inquiries about two of their products. Two types of questions were in each email: some product specific questions (e.g., whether a camera’s frame is metal or plastic), and some service specific questions (e.g., whether they offer gift wrapping or can help write a gift note or send packages without enclosing price information, etc.). We found a large variation in response time (varying from two hours to over five days) across products for the *same* seller (six out of eight sellers provided very different service levels for the two product inquiries). While the above exercise is not conclusive and the different amounts of delay could be due to some random factors, the fact that the same seller takes significantly different amounts of time to respond to similar questions about different products indicates that sellers could indeed be varying their service levels strategically.

warrant direct selling, because the observed not-so-high sales may be due to either a not-so-popular product or a popular product but not-so-good seller services/efforts.

To prevent this to some extent, Amazon requires a baseline level of services from the sellers, and it also expends resources to acquire consumer reviews in an attempt to prevent poor services, which can damage the reputation of its platform. Many anecdotes indicate that Amazon immediately terminates any sellers that are identified as giving poor services. For this reason, sellers always want to provide “acceptable” levels of service to meet Amazon’s standard or normal service levels. However, sellers still have a lot of leeway to decide on how much additional (or “exceptional”) service to provide beyond the standard service level. For example, gift wrapping or other customizable service options, promise of faster shipment, high stock availability and exceptional product support are all beyond the standard service requirements. These factors cannot be costlessly monitored by Amazon but certainly affect the seller’s sales, enabling the seller to mask his high demand from Amazon. Further, the seller’s promotional or other selling efforts that influence the demand but are not directly observed by Amazon are also included in the unobserved services that make demand-masking possible.

While the above discussion is in the context of Amazon, the key forces at play are relevant to online platform retailing in general. Therefore, mid-tail products give rise to an interesting market in which the independent seller benefits from selling on the platform, but he may also be in competition with the platform owner itself. The platform owner can track the seller’s sales to identify whether his product has high-enough demand for it to sell directly. Sales, however, are the outcome of the inherent “popularity level” of the product (due to its design and other attributes) and the seller’s demand-enhancing services, both of which are unobservable to the platform owner. Under the threat of entry by the platform owner, high-demand sellers may attempt to mask themselves as low-demand sellers by providing “acceptable,” but not “exceptional,” service, so that they can continue selling in the future.

This motivates interesting research questions. What implications do such conflicting interests have for the platform owner, the high-demand seller and the low-demand seller? What fee should the platform owner charge? How will different sellers respond in terms of their service provisions? Under

what conditions will the platform owner be able to separate the high-demand mid-tail products from the low-demand mid-tail or long-tail products? Is it ever optimal for the platform owner to forgo its option to sell the product directly? How are consumers affected? Finally, if the platform owner can acquire fully revealing seller reviews, how does this impact the answers to the above questions?

We study the above strategic interactions and provide novel insights into the dynamics of the mid tail of online retailing. First, we find that if the platform owner believes *ex ante* expected demand to be sufficiently high, it will set its fee high enough to separate the high-demand seller from the low-demand seller. In this case, only the high-demand seller will sell on the platform in the early period (separating equilibrium), and the platform owner will subsequently sell the high-demand product directly. In the case of a low *ex ante* probability of high demand, however, the platform owner will set its per-unit fee low enough so that even a low-demand seller will participate on the platform. But this enables the high-demand seller to mimic a low demand seller by under-investing in demand-enhancing services, so that the platform owner will be unable to learn the seller's true demand (pooling equilibrium).

Second, the platform owner may be better off to contractually forgo its option to sell the independent seller's product. This is because, without the threat of entry by the platform owner, the high-demand seller will optimally provide a high level of service and have high sales, from which the platform owner can benefit by charging higher fees. One may expect that sellers prefer less threat of entry. But interestingly, we find that *both* the low-demand and the high-demand seller may *benefit* from the platform owner's threat of entry. This is primarily because if the *ex ante* probability of high demand is small, the platform owner's entry option will lead to a lower per-unit fee than without the threat of entry. Third, the platform owner's entry option can reduce consumer surplus early in the product selling horizon, though it increases consumer surplus late in the selling horizon. Finally, if the platform owner invests in consumer reviews that fully reveal a seller's service level (and, therefore, his true type), then its optimal sales fee will increase (from the no-review case) if the *ex ante* probability of a high-demand type is low, and decrease if that probability is high.

The rest of the chapter is organized as follows. In the next section, we briefly review the related literature. In Section 3.3, we develop an analytical framework to model the interaction between the platform owner and the independent seller. In Section 3.4, we first examine the complete information case; then, we analyze the incomplete information case and compare two scenarios—one, the platform owner credibly commits to not selling the product; two, it retains the option to sell the product in the future. In Section 3.5, we examine the effect of consumer reviews. In Section 3.6, we discuss the robustness of our insights to alternative modeling assumptions. In Section 3.7, we conclude the chapter with a short discussion.

3.2. Review of Relevant Literature

Our work lies at the intersection of internet retailing, platform-based business models, stores within a store, asymmetric information strategies (especially signaling), and signaling under moral hazard. While one and occasionally more than one have been studied, the rich interaction examined here is unique and without much precedent. To clearly delineate our contributions, we briefly discuss the relevant aspects of each literature stream.

Prior work on internet retailing has primarily focused on the interaction between online and offline consumer purchasing (Ansari et al. 2008, Biyalogorsky and Naik 2003, Choi et al. 2010, Choi and Bell 2009, Neslin et al. 2006, Ofek et al. 2009, Pan et al. 2002b, Shankar et al. 2003), the impact of easier online information search on prices (Bakos 1997, Baye et al. 2007, Brynjolfsson and Smith 2000, Pan et al. 2002a), empirically documenting the “long-tail” phenomenon and its implications (Brynjolfsson et al. 2003, 2006, Elberse 2008, Tan and Netessine 2009, Tucker and Zhang 2011), and studying the effects of reviews on firm marketing strategies (e.g., Chen and Xie 2005, Jiang and Srinivasan 2011, Kuksov and Xie 2010). However, as far as we know, our research is the first to identify and analytically study strategic interactions of the aforementioned nature in platform-based internet retailing.

With the advent of new technologies, platform-based business models are becoming increasingly popular. Beyond Amazon’s retail platform, there is a plethora of products and services being turned into a

platform on which sellers and end-users can directly interact and a wide range of products can be offered. Prominent examples include eBay for auctions, iPhone, Android OS and iPad for software applications, and Microsoft Xbox, Sony PlayStation and Nintendo Wii for console-based video games. These developments have motivated the recent literature on two-sided markets (Armstrong 2006, Eisenmann et al. 2006, Parker and Van Alstyne 2005, Rochet and Tirole 2003). This literature primarily focuses on cross-market network effects. In contrast, our focus is not on the platform owner's optimal marketing mix to develop or benefit from its two-sided network. The core of our analysis arises from the aforementioned opposing incentives—strategic learning of demand by the platform owner versus strategic masking by the high-demand seller.

Our work is related to the vast literature on distribution channels in marketing (Coughlan and Wernerfelt 1989, Desai et al. 2004, Desiraju and Moorthy 1997, Iyer and Villas-Boas 2003, Jeuland and Shugan 1983, McGuire and Staelin 1983, Moorthy 1988). Specifically, “stores within a store” (e.g., cosmetics boutiques run by manufacturers in large department stores) can also be considered as platform-based retailing in a physical store. Jerath and Zhang (2010) show that channel efficiency and price competition considerations are the drivers behind the choice of this arrangement. Online platform-based retailing, however, generates a completely different set of issues. First, the number of products sold on online platforms is several orders of magnitude larger than at any physical retailer—which, because of its shelf-space limitation, typically sells only mainstream products—leading to a complex demand-identification problem in our current study. Second, because of large investments and strict, long-term contracts involved, opportunistic behavior on the part of the parent store is limited in a physical store-within-a-store arrangement. However, in the online setting, the platform owner's cherry-picking of third-party sellers' successful products is easily facilitated because of the low investment and the short-term “at-will” nature of the agreement.

Besides contributing to the existing literature on retailing, we also obtain some interesting results for asymmetric information games. First, in most signaling games, a separating equilibrium in which a high-type player separates from a low-type player is the focal equilibrium (e.g., Desai 2000, Desai and

Srinivasan 1995, Moorthy and Srinivasan 1995, Shin 2005, Simester 1995, Soberman 2003). In contrast, in our scenario, a high-demand seller wants to imitate a low-demand seller to avoid the platform owner's entry while a low-demand seller is unconcerned and plays his optimal strategy—the pooling outcome is the focal equilibrium. This is related to the literature on countersignaling (Araujo et al. 2008, Feltovich et al. 2002, Mayzlin and Shin 2009, Teoh and Hwang 1991), but the intent of the high-type player in our case is to hide, rather than reveal, his true type.

Second, most research on signaling does not consider unobservable actions and examines only the signaling of private information from the principal to the agent. In contrast, our research concerns both signaling of private information and unobservable actions. This is similar to Desai and Srinivasan (1995), who study how a franchisor may signal its product's high demand potential to an uninformed franchisee, whose unobservable effort also influences demand. Our model differs structurally in that both the private information about demand and the unobserved effort (service level) are possessed by the same party (the seller) rather than by different parties. More importantly, in our setting, the uninformed party (the platform owner) has to first decide its fee *before* observing the seller's signal about product demand, and subsequently decides whether or not to procure and sell the product directly.

3.3. Model

Consider a new product available for sale on an online retail platform such as Amazon.com. For the ease of understanding and exposition, we will refer to the platform owner as Amazon though our analysis applies to other such retail platforms. Amazon can sell the product directly, or it can let an independent seller offer it and charge him a per-unit fee for each sale. A fixed cost is incurred to sell the new product. Such a fixed cost may include establishing relationships and negotiating contracts with the manufacturers, arranging logistics and allocating warehouse spaces. An independent seller may have a significantly lower fixed cost (for the product under consideration) than Amazon. In fact, the seller's fixed cost may be sunk. For example, the seller may leverage his existing personal connections to procure the product from its manufacturers, and he may use his home basement to store and manage inventory. In addition, the seller

may sell only a few products and thus may not have any costly logistical issues that Amazon faces since it carries hundreds of thousands of products. Collectively, these factors may enable some independent sellers to enjoy a fixed cost substantially lower than that of Amazon. Without loss of generality, we assume that the independent seller has zero or sunk fixed cost whereas Amazon must incur a positive fixed cost ($F > 0$) to sell the new product.

We consider a dynamic model with two time periods. The product can be sold in both time periods and demand in each period i (denoted by a parenthesized superscript $i \in \{1,2\}$) is represented by a linear function: $q^{(i)}(p, e) = \gamma + e^{(i)} - bp^{(i)}$, where p is the price of the product, e is the service (or selling effort) level by the party selling the product (either Amazon or the independent seller), and γ and b are constants.¹⁸ There is uncertainty about the overall product demand—with a prior probability $\theta > 0$, $\gamma = \gamma_H$, and with a probability $1 - \theta$, $\gamma = \gamma_L < \gamma_H$. For ease of exposition, we reframe the uncertainty in the product demand as uncertainty about the seller's type, which is H with probability θ and L with probability $1 - \theta$. The seller knows his type (i.e., whether his product has low or high demand) whereas Amazon knows only the prior probability distribution. This assumption that *ex ante* the seller knows more about the product demand than Amazon is quite reasonable. For example, a local retailer knows which of his products are selling well locally and may decide to sell them online to a larger customer base. An international immigrant may know about a product that sells well in his own country, and can order the product from the manufacturer there to sell on Amazon. In general, it is very plausible that a small seller may have better market demand information than Amazon for the particular product he has identified to sell. Though Amazon may have better knowledge about the demand for many mainstream products, it is reasonable to assume that Amazon does not always have *ex ante* better demand information than *all* third-party sellers for millions of long-tail and niche products. Our interest is, in fact, in those products whose levels of demand are not already known to Amazon.

¹⁸ We assume for analytical tractability that the overall demand intercept is separable in γ and e . However, our insights will hold for other functions that allow γ and e to jointly influence the demand intercept (e.g., a multiplicative form $\gamma^x \cdot (e^{(i)})^y$, $x > 0, y > 0$ will provide the same insights).

We assume that the selling party's service level influences the total demand. A service level e imposes a per-unit marginal cost of $s(e) = ke^2$ on the selling party, where k is a constant and $e \geq 0$. For a more interesting analysis, we assume that the parameters are such that the following conditions hold:

$$\text{C1: (i) } F > \frac{\theta\left(\gamma_H + \frac{1}{4bk} - bc\right)^2 + (1-\theta)\left(\gamma_L + \frac{1}{4bk} - bc\right)^2}{4b} - \frac{[\theta(\gamma_H - \gamma_L) + \gamma_L + \frac{1}{4bk} - bc]^2}{8b}; \quad \text{(ii) } F < \frac{\left(\gamma_H + \frac{1}{4bk} - bc\right)^2}{8b}$$

$$\text{C2: } \gamma_L < \gamma_H \leq \gamma_L + \frac{1}{2bk}.$$

C1(i) ensures that the fixed cost is high enough such that with only the prior information on the seller's type, Amazon will make a higher expected profit by allowing the seller to sell the product than by selling it directly. C1(ii) ensures that the fixed cost F is not too high, such that if Amazon knows that $\gamma = \gamma_H$ (i.e., the demand is high), it prefers selling the product directly rather than letting the independent seller do so. Consequently, when assumption C1 is satisfied, Amazon will allow the independent seller to sell the product if it does not know the true value of γ , but it will sell the product directly if $\gamma = \gamma_H$ and Amazon knows this.

Amazon gains on two grounds by letting the seller offer the product on its platform. First, it earns a fee for each unit sold by the seller. Second, Amazon can acquire more information on market demand since it has direct access to the seller's sales information (price and quantity sold). If the seller's sales are high (i.e., an H-type seller), Amazon will procure and sell the product directly. When Amazon itself sells a product, it can set a lower price than the independent seller because it does not have to pay its own fees (hence avoiding the "double-marginalization" problem). In addition, Amazon can negotiate better prices from the manufacturer, prominently promote its own products to consumers, and offer cheaper/free shipping. Therefore, consumers are much more likely to buy the product from Amazon than from any independent seller if both offer it. As discussed in the introduction, the independent seller's sales will plummet to essentially zero once Amazon directly sells his product. Thus, we make a reasonable

simplification that the independent seller will make zero profit (i.e., Amazon effectively replaces the seller) once Amazon starts directly selling his product.¹⁹

Under Amazon's entry threat, the H-type seller has an incentive to hide his high demand by reducing his services. As discussed earlier, many aspects of service are such that they are transacted outside the platform and are not directly observed by Amazon. Since Amazon observes both the price and quantity sold, essentially the seller's price-quantity pair acts as the seller's signal of his type. The seller can use his unobservable service level to manipulate the signal. Assumption C2 ensures that the H-type seller can use services to create an uninformative signal to prevent Amazon from learning his type. Intuitively, the H-type seller can set his service level and price such that if Amazon observes not-very-high sales, it cannot determine whether this is because $\gamma = \gamma_L$ and the service provided is optimal, or because $\gamma = \gamma_H$ but the service is below the optimal level.

Our model encompasses both asymmetric information (Amazon does not know the seller's type) and moral hazard (Amazon also does not observe the seller's service provision). Previous research in the contracting literature has shown that a menu of contracts can be used by a principal to separate different types of agents under asymmetric information (Lal and Staelin 1986, Rao 1990). Further, a two-part tariff can be used with risk-neutral agents to avoid the moral hazard problem (Holmstrom 1979). We adopt a pure variable fee structure in our model based on the actual contract form adopted by Amazon.²⁰ With millions of products and hundreds of thousands of independent sellers on the platform, a menu of contracts or nonlinear contracts where each contract has multiple components may be very difficult to design or implement. The simplicity of implementing one variable fee structure across all sellers may be the reason why it is adopted by Amazon and many other platforms such as iPhone Apps Store and Google


¹⁹ Our results are robust and qualitatively the same as long as the seller's profit is significantly lower when Amazon sells his product than when Amazon does not sell the product.

²⁰ In addition to its per-unit fee, Amazon does charge the sellers a small fee of \$39 per month, which gives the sellers access to certain services such as adding new products or updating product/seller information displayed. However, since even small professional sellers have monthly sales much higher than many thousands of dollars, this small monthly fee is, in all probability, not levied with the intent of removing moral hazard.

Android Market. Note, however, that using a non-linear contract will not qualitatively change our main results.

The two-period game proceeds as follows. In the first period, “nature” determines the seller’s type. With probability θ the seller is an H-type, and with probability $1 - \theta$ the seller is an L-type. The seller learns his type with certainty whereas Amazon knows only the probability distribution. Amazon first selects a per-unit fee f to charge to the seller. We fix Amazon’s fee to be the same across the two periods because with millions of products sold by independent sellers, a renegotiation or dynamic change of this fee is likely to be costly and is not observed in reality. Given f , the seller chooses whether or not to sell on Amazon; if he decides to sell, he simultaneously chooses his first-period service level $e_t^{(1)}$ and price $p_t^{(1)}$, $t \in \{L, H\}$. Then, the seller’s sales are realized according to the demand $q^{(1)}(p_t^{(1)}, e_t^{(1)}) = \gamma_t + e_t^{(1)} - bp_t^{(1)}$ and both the seller and Amazon realize their respective profits.

Figure 3.2: Extensive form of the game



The extensive form of the game is illustrated in Figure 3.2. At the beginning of the second period, Amazon will update its belief about the seller’s type after observing the seller’s first-period price and

sales. Based on the updated belief, Amazon will decide whether or not to sell the product directly. If Amazon sells the product directly, it will simultaneously choose its service level (e_A) and price (p_A). If Amazon decides not to enter the market, the seller will then simultaneously choose his second-period service level $e_t^{(2)}$ and price $p_t^{(2)}$. Demand is then realized based on the second-period price and service levels. All key notations are summarized in Table 3.2.

Table 3.2: Key Notations

Symbols	Description
t	The independent seller's type, $t = \text{"H"} \text{ or } \text{"L"}.$
θ	The probability (<i>ex ante</i>) that the seller's type is "H."
i	Time period, $i = 1 \text{ or } 2.$
f	The fee Amazon charges the seller for each unit sold.
$e_t^{(i)}$	Type t seller's service level in period $i.$
s	Marginal cost of offering service level e : $s(e) = ke^2$, where $k > 0$ is constant.
$p_t^{(i)}$	Type t seller's price in period $i.$
$q_t^{(i)}$	Type t seller's demand in period i : $q_t^{(i)} = \gamma_t + e_t^{(i)} - bp_t^{(i)}$, where b is constant.
γ_t	Type t seller's demand intercept excluding the effect of his service. $\gamma_H > \gamma_L$
Γ	The overall demand intercept observed by Amazon: $\Gamma \equiv \gamma_t + e_t^{(1)} = q_t^{(1)} + bp_t^{(1)}.$
c	The marginal cost of the product.
F	Fixed cost required for Amazon itself to sell a product directly.
$\mu(H p_t^{(1)}, q_t^{(1)})$	Amazon's posterior belief that the seller is of type H after observing $p_t^{(1)}, q_t^{(1)}.$
$\Pi_A^{(i)}$	Amazon's expected profit in period $i.$
Π_A	Amazon's overall expected profit for both periods $i = 1 \text{ or } 2.$
$\Pi_t^{(i)}$	Type t seller's profit in period $i.$
Π_t	Type t seller's overall profit for both periods $i = 1 \text{ or } 2.$
$\hat{}$	The hat over a variable indicates the case of completion information.
$\bar{}$	The bar over a variable indicates the case of no entry threat by Amazon.
"sep", "pool"	These subscripts indicate the separating and pooling outcome, respectively.
"rev"	This subscript indicates the variable is for the case with seller reviews.

3.4. Analysis

We organize our analysis in the following way. In Section 3.4.1, we analyze a benchmark case of complete information, in which Amazon knows the true demand (i.e., the seller's type). In Section 3.4.2, we analyze another benchmark case of asymmetric information and unobservable service, in which Amazon contractually commits to not selling the product. In Section 3.4.3, we examine our focal case of asymmetric information and unobservable service in which Amazon retains the option to sell the product.

3.4.1 Complete Information

In this section, we examine a complete-information game in which Amazon knows the seller's type. Without information asymmetry, the game becomes simple to solve. If the demand is low ($\gamma = \gamma_L$), Amazon will let the independent seller sell it in both periods. Given Amazon's fee \hat{f} , the L-type seller will select the same service levels and prices for both periods $i \in \{1, 2\}$. (We will use a "hat" over the variable to indicate the case of complete information.) The L-type seller's total profit from both periods is given by $\hat{\Pi}_L = \sum_{i=1}^2 \left\{ \left(\gamma_L + \hat{e}_L^{(i)} - b\hat{p}_L^{(i)} \right) \left[\hat{p}_L^{(i)} - c - \hat{f} - k(\hat{e}_L^{(i)})^2 \right] \right\}$, where $c > 0$ is the product's marginal cost.²¹ Solving the first order conditions, one easily finds the L-type seller's optimal (first-best)

service level and price: $\hat{e}_L^{*(i)} = \frac{1}{2bk}$ and $\hat{p}_L^{*(i)} = \frac{\gamma_L + b(c + \hat{f}) + \frac{3}{4bk}}{2b}$. The corresponding profit is given by

$\hat{\Pi}_L^*(\hat{f}) = \frac{\left[\gamma_L - b(c + \hat{f}) + \frac{1}{4bk} \right]^2}{2b}$. Amazon's profit in the case of an L-type seller is given by $\hat{\Pi}_{A,L}(\hat{f}) =$

$\sum_{i=1}^2 \left\{ \hat{f} \hat{q}_L^{(i)}(\hat{f}) \right\}$, where $\hat{q}_L^{(i)}(\hat{f}) = \gamma_L + \hat{e}_L^{*(i)} - b\hat{p}_L^{*(i)}(\hat{f})$ is the seller's quantity sold in period i as a

function of \hat{f} . Thus, Amazon's optimal fee is $\hat{f}^* = \frac{\gamma_L - bc + \frac{1}{4bk}}{2b}$ and its profit from the L-type seller is

$$\hat{\Pi}_{A,L}^* = \frac{\left(\gamma_L - bc + \frac{1}{4bk} \right)^2}{4b}.$$

²¹ For analytical simplicity and without loss of generality, we assume the discount rate to be one across the two periods. Our main results stay qualitatively the same if we allow for discounting of the second-period profit; only the parameter regions of those results will change depending on the discount rate.

If the demand is high ($\gamma = \gamma_H$), Amazon will sell the product directly in both periods. In this case, its optimal service level and price are easily found to be $e_A^{*(i)} = \frac{1}{2bk}$ and $p_A^{*(i)} = \frac{\gamma_H + bc + \frac{3}{4bk}}{2b}$, with a corresponding total profit of $\hat{\Pi}_{A,H}^* = 2 \left[\frac{(\gamma_H - bc + \frac{1}{4bk})^2}{4b} - F \right]$.

3.4.2 The Platform Owner Commits to “No Entry” in the Second Period

In this section, we analyze the case of asymmetric information and unobservable service in which Amazon contractually commits to not selling the product in the future.²² This removes the H-type seller’s incentive to mask his demand. Therefore, with Amazon’s credible commitment to “no entry,” the seller, irrespective of his type, will choose his “first-best” service level and price according to his type, given Amazon’s per-unit fee. We will later analyze whether or not it is advantageous for Amazon to forego its option of entry.

The seller will select the same service levels and prices for both periods $i \in \{1, 2\}$. His total profit from both periods is given by $\bar{\Pi}_t = \sum_{i=1}^2 \left\{ \left(\gamma_t + \bar{e}_t^{(i)} - b\bar{p}_t^{(i)} \right) \left[\bar{p}_t^{(i)} - c - \bar{f} - k(\bar{e}_t^{(i)})^2 \right] \right\}$, where $t \in \{L, H\}$ represents the seller’s type. Given Amazon’s per-unit fee \bar{f} , the seller chooses the service levels ($\bar{e}_t^{(i)}$) and prices ($\bar{p}_t^{(i)}$) to maximize his total profit. (Throughout the chapter, we use a “bar” over any variable to indicate that the variable is for the case of no entry threat by Amazon.)

Lemma 1: *Without threat of entry by the platform owner, both types of sellers will offer the same optimal service levels in both periods; their prices and profits differ (separate) according to their types.*

Lemma 1 shows that if Amazon commits to not entering the market, both types of sellers will, in equilibrium, offer the same (high) service level. The prices and profits, however, differ across the two types.²³ The equilibrium outcome (service levels, prices, fee, and overall profits for a seller of type t and for Amazon) are provided in the Appendix.

²² Amazon can make a legally-binding credible commitment of this nature by putting in its seller agreement something like: “Amazon.com Inc. cannot contract with a product’s original manufacturer to directly sell the product in competition with the third-party seller without specific permission from the seller.”

²³ All proofs in this paper are relegated to the Appendix.

3.4.3 With Threat of Entry by the Platform Owner

In this section, we study our focal case in which Amazon keeps its option to sell the product directly in the second period and will do so if it identifies the seller as an H-type seller. Under different parameter conditions, we obtain either a separating equilibrium (in which Amazon can determine the seller's type after the first period) or a pooling equilibrium (in which it cannot).

The Platform Owner Learns the Independent Seller's Type: Separating Equilibrium

In this equilibrium, Amazon is able to learn the seller's true type after the first period. Amazon will directly sell the product in the second period only if it identifies the seller as an H-type. Note that it is not possible to have a separating equilibrium with any fee $f < \frac{\gamma_L - bc + \frac{1}{4bk}}{b}$. This is because at such a fee both types of sellers sell on the platform and the H-type strictly prefers mimicking the L-type to avoid Amazon's entry. (Given f , the H-type's separating equilibrium profit is weakly less than what he can earn in two periods if he mimics the L-type in the first period to deter Amazon's entry). In contrast, any fee f satisfying $\frac{\gamma_L - bc + \frac{1}{4bk}}{b} \leq f < \frac{\gamma_H - bc + \frac{1}{4bk}}{b}$ will induce a separating equilibrium since with such a fee, only the H-type seller will profitably sell a positive quantity (the L-type will not enter the market).²⁴ Thus, the H-type seller does not have any incentive to mimic the L-type and will choose his first-best service level and price (conditional on f). In other words, the H-type seller makes a positive profit in the first period and zero profit in the second period, in which Amazon will sell the product directly.²⁵ Therefore, with probability θ , Amazon earns some fees from the H-type seller in the first period, and will sell the product itself in the second period. If the seller is an L-type, which happens with probability $1 - \theta$, Amazon will earn zero profit. We use the subscript "sep" to indicate the separating outcome. In the Appendix, we

²⁴ The upper bound is needed because with $f \geq \frac{\gamma_H - bc + \frac{1}{4bk}}{b}$, neither type of seller can profitably sell on Amazon.

²⁵ According to his demand function $q_L = \gamma_L + e_L - bp_L$, to sell any positive quantity, the L-type seller must charge a price $p_L < \frac{\gamma_L + e_L}{b}$. If $f \geq \frac{\gamma_L - bc + \frac{1}{4bk}}{b}$, the L-type's profit margin becomes $p_L - c - f - ke_L^2 < \frac{\gamma_L + e_L}{b} - c - \frac{\gamma_L - bc + \frac{1}{4bk}}{b} - ke_L^2 = \frac{e_L - bke_L^2 + \frac{1}{4bk}}{b} = -k \left(e_L - \frac{1}{2bk} \right)^2 < 0$. Hence, he will not sell on Amazon if its fee is so high.

compute f_{sep}^* , $\Pi_{t,sep}^*$, and $\Pi_{A,sep}^*$. Later, we specify in Proposition 1 the condition under which this separating equilibrium is realized.

The Platform Owner Does Not Learn the Independent Seller's Type: Pooling Equilibrium

We now consider the case of $f < \frac{\gamma_L - bc + \frac{1}{4bk}}{b}$, in which both types of sellers will sell on Amazon's platform.

Note that if the realized equilibrium corresponds to $f < \frac{\gamma_L - bc + \frac{1}{4bk}}{b}$, then a pooling equilibrium arises because given such a fee, the H-type seller strictly prefers mimicking the L-type to avoid Amazon's entry in the second period. Since the seller knows his own type with certainty, for sequential equilibrium, we need to specify only Amazon's posterior belief about the seller's type. There can be infinitely many pooling equilibria supported by different out-of-equilibrium beliefs. However, in our case, all such equilibria except one are ruled out or refined away by specifying what beliefs are “unreasonable” using the “intuitive criterion” by Cho and Kreps (1987) and additional logical reasoning. In the Technical Appendix, we show that such refinements lead to the following unique pooling equilibrium: In the first period, both types of sellers choose $p_t^{*(1)}(f) = \frac{\gamma_L + b(c+f) + \frac{3}{4bk}}{2b}$ and sell a quantity of $q_t^{*(1)}(f) = \frac{\gamma_L - b(c+f) + \frac{1}{4bk}}{2}$. This outcome corresponds to the L-type seller choosing his first-best service level and price, given f . Below is Amazon's posterior belief that supports this unique pooling equilibrium. For convenience, we define the overall demand intercept as $\Gamma \equiv q_t^{(1)} + bp_t^{(1)} = \gamma_t + e_t^{(1)}$.

$$\mu(H|p_t^{(1)}, q_t^{(1)}) = \begin{cases} 1, & \text{if } \Gamma > \gamma_L + \frac{1}{2bk}, \\ \theta, & \text{if } \gamma_H \leq \Gamma \leq \gamma_L + \frac{1}{2bk}, \\ 0, & \text{if } \Gamma < \gamma_H \end{cases} \quad (1)$$

Again, we remind the reader that Proposition 1 will specify the condition for this pooling equilibrium outcome to occur. We now derive this equilibrium.

In the second (and last) period, the seller faces no future entry threat and will thus choose his first-best service level and price, conditional on Amazon's fee f . Hence, in the second period, a seller of

type t will choose $e_t^{*(2)} = \frac{1}{2bk}$ and $p_t^{*(2)}(f) = \frac{\gamma_t + b(c+f) + \frac{3}{4bk}}{2b}$, which are the same as (A3) and (A5), respectively. The seller's second-period profit is given by $\Pi_t^{*(2)}(f) = \frac{[\gamma_t - b(c+f) + \frac{1}{4bk}]^2}{4b}$.

Amazon observes the seller's first-period price $p_t^{(1)}$ and unit sales $q_t^{(1)}$ since it processes all orders on its platform. Effectively, $p_t^{(1)}$ and $q_t^{(1)}$ are a multi-dimensional signal of the seller's type. After the first period, having observed $p_t^{(1)}$ and $q_t^{(1)}$, Amazon learns the overall demand intercept $\Gamma = \gamma_t + e_t^{(1)} = q_t^{(1)} + bp_t^{(1)}$, but it may not be able to deduce γ_t or the seller's type since it does not observe $e_t^{(1)}$.

Note that an L-type seller has no incentive to prevent Amazon from learning his true type, because with Amazon knowing that he is an L-type, it will not enter the market. Note also that at any given f that is not too high to preclude the L-type seller from selling profitably, the L-type seller is actually indifferent between the pooling and separating outcomes since both outcomes lead to no entry by Amazon. This implies that an L-type seller will play his first-best strategy as long as Amazon's belief will not falsely identify him as an H-type for doing so, which is the case for the belief system (1). In contrast, an H-type seller facing the threat of entry by Amazon has an incentive to strategically choose his first-period service level and price to exactly mimic the L-type seller's first-best price and sales. Such a strategy by the H-type seller results in a unique pooling outcome and prevents Amazon from learning his true type. Therefore, Amazon will not enter the market in the second period.

We now analyze the implications of this pooling equilibrium in detail. At this equilibrium, Amazon is unable to determine the seller's type after the first period and hence will not directly sell the product in the second period. Since the second period is the final period, the seller will face no future entry threat by Amazon and hence will set the service level and price to maximize his profit according to his true demand. Thus, given f , the seller's optimal service level and price in the second period are the same as those in Section 3.4.2. We use the subscript "pool" to indicate that a variable corresponds to the pooling outcome. The possible pooling equilibrium outcome is provided in the Appendix.

Realized Equilibrium

Proposition 1 shows that when the probability of H-type is below a threshold, the pooling equilibrium is realized, otherwise the separating equilibrium is realized.

Proposition 1: For any set of values of the other parameters, there exists $\theta^* \in (0, 1)$ such that the pooling equilibrium outcome is realized if $\theta < \theta^*$ and the separating equilibrium outcome is realized if $\theta \geq \theta^*$.

Figure 3.3: Market outcomes in (γ_H, θ) parameter space

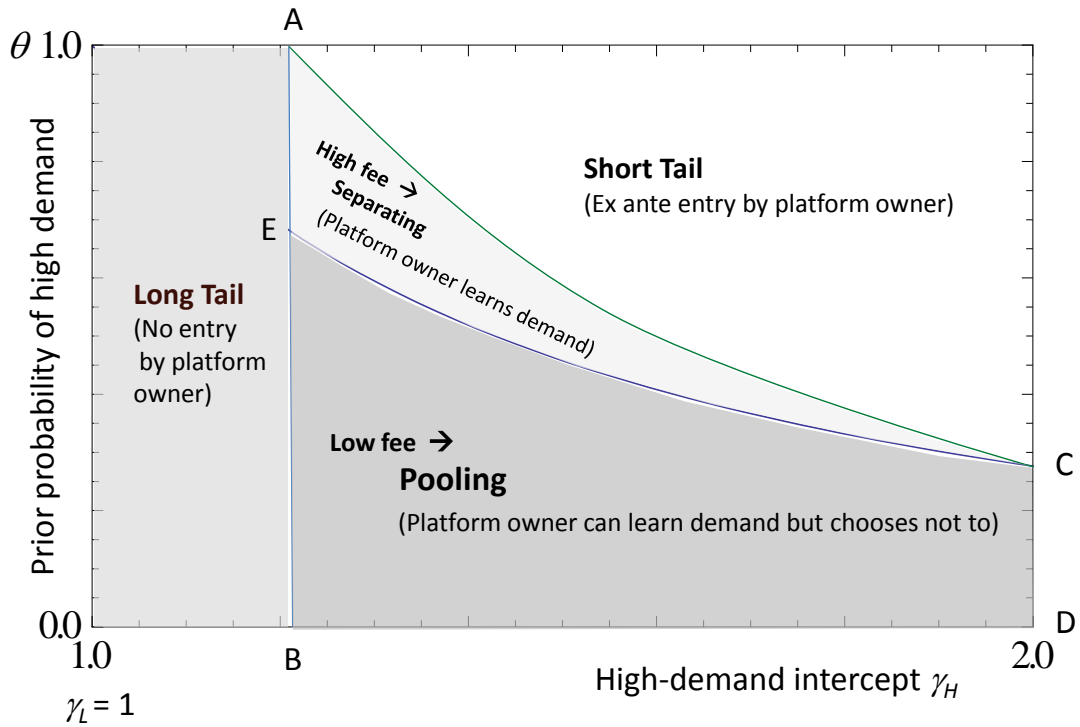


Figure 3.3 illustrates the equilibrium realizations in the (γ_H, θ) parameter space. In the figure, the curve AC corresponds to the boundary defined by assumption C1(i); the line AB corresponds to the boundary defined by assumption C1(ii); the line CD corresponds to the right-side boundary of assumption C2; the curve EC corresponds to θ^* . If the *ex ante* expected demand is high enough (i.e., θ and γ_H are in the “Short Tail” region), Amazon will have entered the market in the *first* period. Products such as a new version of a highly popular digital camera likely have parameters that fall into this region; there is still demand uncertainty for such products, but their expected demand is high enough to warrant Amazon’s

direct selling immediately. The left rectangular “Long Tail” region represents the very low-demand, long-tail products for which even γ_H is small enough such that Amazon will not be able recover its fixed cost if it sells them directly. Our analysis has focused on the most strategically-interesting parameter region—the mid-tail region where the expected demand is in the middle range. We find two qualitatively different mid-tail regions. In the “Separating” parameter region, the expected demand is relatively high and at equilibrium Amazon will set its fee high enough to separate both types of sellers to directly sell the high-demand product in the second period. In the “Pooling” region, the expected demand is relatively low, Amazon actually finds it optimal not to learn the seller’s true type, and will set a fee low enough to allow both types of sellers to enter the market. In that case, in the first period, the H-type seller will mimic the L-type by strategically lowering his service level, and in the second period Amazon will not enter the market since it cannot learn the seller’s type. In contrast to the extant literature on signaling which has focused on separating equilibrium outcomes, the focal outcome in our research is the pooling equilibrium.

Now, we examine how Amazon’s entry option impacts its own profit.

Proposition 2: *If $\theta < \theta^*$, the platform owner’s fee and its expected profit are both lower if it retains its entry option than if it forgoes it.*

Proposition 2 shows that in the pooling parameter region, Amazon’s entry option will hurt its own profit. Amazon is worse off in the pooling equilibrium than if it forgoes its future option to sell the product directly. This is because Amazon’s threat of entry gives the H-type seller an incentive to reduce his first-period service level to mimic the L-type seller so that Amazon is unable to learn his type and hence will not enter the market in the second period. If Amazon *ex ante* forgoes its entry option, the H-type seller will then optimally provide a high service level even in the first period, and in turn, Amazon will have an incentive to charge a higher fee to benefit from the H-type seller’s high sales. As a result, in the pooling parameter region, Amazon’s fee is lower if it retains an entry option than if it forgoes it. Amazon’s overall profit is also lower when it retains the entry option because of both its lower per-unit fee and the H-type seller’s strategically reduced first-period sales.

One managerial implication is that if Amazon can do so costlessly, it should commit to no-entry for products with parameters that fall in the pooling region. In addition, our analysis shows that Amazon's optimal fee may differ across products depending on the parameters associated with the products. However, in practice, with tens of millions of products on its platform, Amazon cannot very cost-efficiently contract on a product-by-product level; in fact, its fees vary only across product categories, e.g., 6% for computers, 8% for cameras, 10% for tires and wheels, 12% for musical instruments, 15% for toys and video games, etc. Given that in practice Amazon uses only one blanket seller-contract, if it commits to no-entry, it solves the moral hazard problem, but it will give up the profit potential from cherry-picking of the third-party sellers' high-demand products. According to Amazon's annual reports, Amazon makes most of its profits from direct selling. Even though it directly sells only about 7% of the products listed on its platform, its sales accounts for 69% of all unit sales on its platform. Thus, it is understandable that if Amazon uses a blanket contract for all third-party sellers, it does not want to forgo its entry option.

Now we examine how Amazon's entry option affects the third-party seller and the consumers. Intuitively, one may expect that a seller prefers less threat of entry and that consumers will benefit from potential competition in the market. However, we find that this is not necessarily the case.

Proposition 3: *If $\theta < \theta^*$, the L-type seller makes a higher profit with the platform owner's threat of entry than without it. The H-type seller makes a higher profit with the platform owner's threat of entry than without it only if $\gamma_H < \gamma^*$. (Both θ^* and γ^* are constant expressions given in the appendix.)*

According to Proposition 3, in the pooling parameter region, the L-type seller will benefit from Amazon's threat of entry. The intuition is that with or without Amazon's entry threat, the L-type seller will choose the same service level and sets his first-best price given Amazon's fee, but Amazon's fee is lower when it retains its entry option. Of course, if the separating outcome is realized as in the case of $\theta > \theta^*$, the L-type seller is hurt by Amazon's entry option because he sells nothing on the platform.

More surprisingly, the H-type seller can also benefit from Amazon's threat of entry; this happens if γ_H is not too large. With no entry threat, the H-type's first-period service is higher since he need not

mask his demand, but Amazon's fee is also higher because it now has incentives to raise the fee to benefit from the H-type's high demand. In contrast, when Amazon retains an entry option, the H-type will lose out on unit sales in the first period as he mimics the L-type to prevent Amazon's entry. However, all other factors benefit him including a lower fee for both periods and a lower service cost in the first period. Intuitively, if γ_H is not too large relative to γ_L , the H-type seller's forgone profit in the first period due to his strategic reduction of sales will be more than compensated for by the gain in the second-period profit from his high sales at lowered fees. But if γ_H is very high, then the H-type seller's lost unit sales in the first period will dominate, yielding a lower overall profit than if Amazon commits to "no entry."

Proposition 4: *If $\theta < \theta^*$, the platform owner's threat of entry increases the second-period consumer surplus. Furthermore, it increases the first-period consumer surplus in the case of an L-type seller and decreases the first-period consumer surplus in the case of an H-type seller.*

Proposition 4 shows the effect of Amazon's threat of entry on the overall consumer surplus when the probability of an H-type is relatively low (in the pooling parameter region). The consumer surplus in the second period is clearly higher when Amazon keeps its entry option than not. This is because in the second period both types of sellers choose their first-best service levels and prices given Amazon's fee (as in the case of no entry threat) and this fee is lower when Amazon retains its entry option, which leads to lower prices to the consumer. For the same reason, in the case of an L-type seller, the first-period consumer surplus is also higher when Amazon keeps its entry option. But interestingly, in the case of an H-type seller, the first-period consumer surplus is actually lower with Amazon's threat of entry than without it (in the pooling parameter region). This is because even though Amazon's entry option leads to a lower price, it induces the H-type seller to provide a significantly lower service level in the first period to reduce his sales to that of an L-type. That is, in the first period because of the H-type seller's lowered services, significantly fewer consumers will buy from him than if Amazon had committed to "no entry," in which case the H-type seller would provide a high service level to benefit from his full demand

potential. The net effect is that in the case of an H-type seller, Amazon's threat of entry reduces the first-period consumer surplus due to the H-type seller's reduced services to mask his demand.

Proposition 5: (a) *The steeper the demand curve (the larger b), the lower the platform owner's fee.*

(b) *The platform owner's fee increases with θ when $\theta < \theta^*$ and is independent of θ when $\theta > \theta^*$.*

We obtain the intuitive result that a steeper (i.e., more elastic) demand leads to a lower fee by Amazon. A more elastic demand intuitively means that for any given decrease in the price to the consumer, the quantity demanded will increase more. Thus, when demand is more elastic, Amazon tends to have more incentive to reduce its fee so as to get the seller to reduce his prices to sell many more units, which will lead to higher total fees. We find that Amazon's pooling equilibrium fee increases as the probability of the seller being an H-type increases. But as this probability becomes large enough, Amazon will effectively be targeting only the H-type seller (to obtain the separating outcome) and its fee is thereby set optimally conditional on the seller being an H-type. Therefore, for large θ , Amazon's fee will become constant. In the low-probability event of an L-type seller, Amazon makes zero profit, but its expected profit is maximized with the high fee, which induces a separating equilibrium.

3.5. Effects of Consumer Reviews

Though Amazon may not directly observe the seller's service level, it can solicit reviews from the seller's first-period customers (since Amazon has their contact information). Such reviews can, to some extent, help Amazon estimate the seller's service level. Once Amazon knows the service level, it will be able to better infer the inherent product demand.²⁶

Here we assume the extreme case that consumer reviews fully reveal the seller's service level. After the first period, Amazon will acquire such reviews to determine the seller's service level and hence correctly infer the seller's type. Note that the fully-revealing reviews do not imply a full (complete) information game, because Amazon will know the seller's type only after the first period.

²⁶ To some extent, Amazon may also use product reviews to infer the demand. Thus, in our model, such product reviews serve a similar role to seller reviews in that they help Amazon to learn the seller's type after the first period.

Since the seller knows that consumer reviews will reveal his service level and allow Amazon to learn his type, he no longer has any incentive to deviate from his first-best service and price levels as given in (A3) and (A5). After the first period, Amazon will learn the seller's true type via his reviews and will directly sell the product if the H-type seller is revealed. In the case of an L-type seller, Amazon will let him continue in the second period. Given Amazon's fee f_{rev} , both types of sellers will choose

$$e_t^* = \frac{1}{2bk} \text{ and } p_t^*(f_{rev}) = \frac{\gamma_t + b(c + f_{rev}) + \frac{3}{4bk}}{2b}, \text{ which yields a first-period profit of } \Pi_t^{*(1)} = \frac{[\gamma - b(c + f_{rev}) + \frac{1}{4bk}]^2}{4b}.$$

In the case of an H-type seller, Amazon's first-period profit is $f_{rev} \cdot \frac{\gamma_H - b(c + f_{rev}) + \frac{1}{4bk}}{2}$. In the second period, Amazon will sell the product directly, making a maximum profit of $\frac{(\gamma_H - bc + \frac{1}{4bk})^2}{4b} - F$ at the optimal price of $p_A^* = \frac{\gamma_H + bc + \frac{3}{4bk}}{2b}$. With the L-type seller, Amazon makes the same profit of $f_{rev} \cdot \frac{\gamma_L - b(c + f_{rev}) + \frac{1}{4bk}}{2}$ for each period. Therefore, Amazon's expected total profit for both periods is given by:²⁷

$$\Pi_{A,rev}(f_{rev}) = \theta \left[f_{rev} \cdot \frac{\gamma_H - b(c + f_{rev}) + \frac{1}{4bk}}{2} + \frac{(\gamma_H - bc + \frac{1}{4bk})^2}{4b} - F \right] + 2(1 - \theta) f_{rev} \cdot \frac{\gamma_L - b(c + f_{rev}) + \frac{1}{4bk}}{2}.$$

It is straightforward to show that Amazon's optimal fee and profit are respectively given by

$$f_{rev}^* = \frac{\theta(\gamma_H - bc + \frac{1}{4bk}) + 2(1 - \theta)(\gamma_L - bc + \frac{1}{4bk})}{2b(2 - \theta)} \quad \text{and}$$


$$\Pi_{A,rev}^* = \frac{[\theta(\gamma_H - bc + \frac{1}{4bk}) + 2(1 - \theta)(\gamma_L - bc + \frac{1}{4bk})]^2}{8b(2 - \theta)} + \theta \left[\frac{(\gamma_H - bc + \frac{1}{4bk})^2}{4b} - F \right].$$

Proposition 6: *With fully revealing consumer reviews, the platform owner's fee is higher when $\theta < \theta^*$ and lower when $\theta > \theta^*$ (compared with the case of no reviews).*

²⁷ We have assumed that $\theta < \frac{2(\gamma_L - bc + \frac{1}{4bk})}{\gamma_H - bc + \frac{1}{4bk}}$, i.e., θ is small enough such that both types of sellers will be targeted by

Amazon. Otherwise, Amazon's optimal decision will be to completely ignore the possibility of the L-type seller.

Figure 3.4: Effect of reviews on Amazon's fee



Proposition 6, illustrated in Figure 3.4, shows that with fully revealing reviews, Amazon's fee will be higher when θ is small (i.e., in the pooling parameter region) and lower when θ is large (i.e., in the separating parameter region). Recall that without reviews, if θ is small, Amazon has an incentive to lower its fee from the no-entry-threat level since the H-type seller will mimic the L-type in the first period. With fully revealing reviews, the H-type will not be able to mask his demand, and will thus choose his first-best service and price levels and sell a high volume in the first period, which gives Amazon the incentive to charge a higher fee. In contrast, if θ is large (in the separating equilibrium) and there are no reviews, Amazon loses the potential profit from the L-type seller. With fully revealing reviews, Amazon no longer needs to set its fee high to separate the two types of sellers, because reviews will reveal the seller's true type. Thus, Amazon will optimally reduce its fee to capture some profits from the L-type seller.

Not surprisingly, reviews make the H-type seller worse off since he will make zero profit in the second period when Amazon enters the market. If θ is small, consumer reviews will also reduce the L-type seller's profit because of Amazon's increased fee. But when θ is large, consumer reviews will make the L-type seller better off because Amazon's reduced fee now allows him to sell profitably (whereas without reviews he will not be able to sell profitably). Amazon's incentive to acquire consumer reviews can be represented by its potential gain in profit: $\Delta \equiv \Pi_{A,rev}^* - \Pi_A^*$. We find that in the pooling parameter region, the larger θ is, the more incentive Amazon has to invest in reviews. In contrast, in the separating

parameter region, Amazon already learns the seller's type with its optimal high fee; its expected profit gain from reviews comes from the L-type seller, whose likelihood decreases as θ increases. Hence, in the high θ range, Amazon's incentive to acquire reviews will decrease with θ .

3.6. Robustness to Alternative Assumptions

In this section, we discuss the robustness of our insights to several alternative modeling assumptions. First, our model has focused on the uncertainty in the demand intercept (γ). Alternatively, the uncertainty in demand may come from the slope of the demand curve (b) rather than the intercept. In such an alternative model, the H-type seller corresponds to the less elastic demand (i.e., smaller b) and the L-type seller corresponds to the more elastic demand (i.e. larger b). The analysis for such a model is very similar to the one that we have done; our key insights remain qualitatively the same.

Second, in practice, platform owners typically charge the sellers a per-unit fee based on the proportion of their sales price; for example, Amazon's fee ranges from 6% to 25% of the sales price depending on the product category. We have assumed that the per-unit fee is fixed (not dependent on price). A fee proportional to price will clearly influence the seller's pricing decisions. However, this does not qualitatively alter the strategic tradeoffs between the platform owner and the seller, and hence our key results remain qualitatively the same (though the parameter region for the pooling outcome will be different). In particular, the seller will have more incentives to charge a lower price because, intuitively, with a proportional fee, a lower price leads to a lower fee and a lower overall marginal cost to the seller (which has three components: the wholesale price, the fee paid to Amazon, and the service cost). We find through numerical examples that this intuition is indeed correct. Since our focus is on the strategic interactions arising from information asymmetric and moral hazard, rather than on the optimal price per se, using such a proportional fee only leads to more analytical complexity while it does not yield additional insights into our research questions.

Third, we have assumed the same demand for both periods. If the product demand varies across periods, we expect our results to qualitatively hold as long as there is a large-enough positive correlation

between the demands across the periods. In particular, if the second-period demand intercept is a multiple α of the first-period demand, then the similar analysis will carry through—both the platform owner and the seller need to adjust the demand by the factor α (which can be larger or small than one) in the second period. While the analysis is more complex with the addition of a new parameter, the strategic considerations of both the platform owner and the seller remain qualitatively the same. For example, if demand falls significantly over time, Amazon will be more likely to not enter the market or exit the market to encourage entry by the third-party sellers.

Fourth, though our model has two time periods, it can be generalized to any finite number of time periods. Our main pooling outcome, for example, will become the following. For a number of initial periods, the H-type seller will mimic the L-type seller; after that, he will stop mimicking and choose his first-best service and price for all later periods. Essentially, when there are only a small number of periods left, the platform owner will no longer find it worthwhile to enter the market even if it identifies the H-type seller at that point because the product will have reached the later stage of the product life cycle or because a new version of the product will soon be released.

Fifth, in our model, we use a pure variable fee rather than a non-linear fee that depends on the quantity sold by the seller (perhaps with a quantity discount). A pure variable fee corresponds to reality in most platform-based retailing settings. However, our main results regarding the pooling equilibrium remain qualitatively the same even if we adopt a non-linear fee. In the Technical Appendix, we show that our focal pooling equilibrium still exists in certain parameter regions even when a general non-linear fee structure is used.

Sixth, we have assumed that Amazon's fee does not change across periods. If a dynamic fee is allowed in our model, Amazon may make more profit without entering the market in the second period after identifying the high-type seller (and thus avoid incurring the entry fee). In theory, after identifying the high-type seller, Amazon can change its second-period fee to extract essentially all the profit from the seller, and make more profit than if it enters the market itself. This separating outcome, though appearing different, is actually the same from the high-type seller's perspective. The high-type seller makes zero

profit in the second period if he is replaced by Amazon or if Amazon optimally changes its fee to extract away his entire surplus. In the case of our focal pooling outcome, allowing for a dynamic fee in our model will not give Amazon any reason to change its fee in the second period since Amazon has the same information about the seller's type at the beginning of the first period and the beginning of the second period (the same as the prior probability).

Finally, we have implicitly normalized the seller's outside option to have a utility of zero rather than explicitly modeling any dynamic outside option of the seller. However, our main intuition and results remain qualitatively the same if we explicitly introduce a parameter to represent the seller's utility for his outside option. Intuitively, we expect that as the seller's outside option becomes more attractive, Amazon will have more incentive to stay out of the market to attract entry by the sellers since they will otherwise be more likely to not enter the platform.

3.7. Conclusion

As online retailing continues to grow, major retailers such as Amazon.com are relying heavily on the platform model of selling. Many small independent sellers utilize the retailing platform to sell products not carried directly by the platform owner. Platform retailing is a win-win for all—consumers get easy access to their preferred but rare products, the small companies get access to these consumers, and the platform owner keeps a percentage of the independent sellers' revenues. However, anecdotal evidence suggests that there is an interesting dynamic prevalent in platform retailing. The platform owner has an incentive to let the independent sellers offer products on its platform, observe the sales of these products, and then cherry-pick the products with high sales potential to procure and sell directly, effectively driving the independent seller's sales to zero. Anticipating the platform owner's demand-learning and cherry-picking incentives, the independent seller has an incentive to hide any high demand by strategically lowering his services to reduce his early sales to prevent the platform owner from learning the true demand. The platform owner, in turn, needs to decide how to set its fee knowing a high-demand seller's incentive to mask his demand.

As an outcome of these strategic interactions, we find that even though the platform owner can set a high enough fee to identify the high-demand seller, it may not always be optimal to do so. If the probability of the seller being an H-type is relatively low, the platform owner will charge a low fee such that both types of sellers sell on the platform. This results in a pooling equilibrium in which the platform owner is unable to learn the true demand because the H-type seller can mask his demand by under-investing in services. Due to this inefficiency, the platform owner is, surprisingly, worse off by keeping its option of entry. Furthermore, the platform owner's threat of entry may benefit both types of sellers because they have to pay a lower fee, while it may reduce or increase the consumer surplus.

Service standardization and monitoring by the platform owner tend to reduce but not completely remove the H-type seller's ability to use a reduced services or efforts to mask his demand. The platform owner can also invest in acquiring consumer reviews. With fully revealing consumer reviews, the platform owner will be able to learn the true demand after the initial time period. We find that such consumer reviews benefit the platform owner and will hurt both types of sellers in the pooling parameter region (with a low *ex ante* probability of the seller being an H-type). But even good consumer reviews are unlikely to fully reveal the seller's service level, because such reviews tend to reveal post-sale service levels rather than pre-sale services. With the seller's mediocre pre-sale service, some consumers may not buy the product and hence cannot write seller reviews as is the case on Amazon. Thus, to the extent that consumer reviews may not fully reveal the seller's service level, the H-type seller may still be able to strategically mask his demand.

Interestingly, our framework can be reinterpreted to provide insights into non-platform retailing situations as well. Suppose that a manufacturer introduces a new product in a certain market. The manufacturer is not certain whether the product will have a high demand (γ_H) or a low demand (γ_L), but has a prior probability of θ for the high demand. The manufacturer may sell through a local retailer who can privately observe the demand potential (γ_H or γ_L) and decide how much promotional effort (e) to invest. The manufacturer collects a wholesale price (f) from the local retailer, but has incentives to go

direct if it learns the product demand is high but not if the demand is low. Such a setting is isomorphic to our model. Our analysis suggests that unless the manufacturer has a high-enough prior for a hit product, it should commit to the local retailer that it will not go direct.

Our study is the first to explore the strategic interactions between the platform owner and the independent sellers in the mid tail of online platform-based retailing, and offers several avenues for future research. For instance, motivated by the actual contract structure on Amazon, we have assumed a single per-unit fee contract. With hundreds of thousands of independent sellers, even if the platform owner uses a menu with several options to distinguish between different sellers and classify them, our analysis will be relevant for the thousands of independent sellers within each class. Nevertheless, extending our framework to a menu of contracts may yield interesting insights with regard to separating outcomes.

We analyze a case with one third-party seller (of either H or L type) and a monopoly platform owner. However, there could be competition at both levels. If multiple sellers on the platform sell the same product, there can still be a symmetric pooling equilibrium in which all sellers prefer to mask the high demand. However, the incentive for each seller to mask the high demand will reduce because of possible free riding from other sellers. As the number of sellers increases, we expect that the free riding problem will become more severe and that the symmetric, pooling outcome may become less likely to be the realized equilibrium outcome. If we introduce competing platforms in our model, our main results and intuitions will hold as long as each platform has a segment of loyal online consumers. Furthermore, with competing platforms, the unit fees charged by the platform owners will tend to be lower because of competition. With lower fees, it is more likely that sellers of all demand types will enter the platforms. Hence, intuitively, we expect that the existence of competing platforms will make the pooling outcome (our focal equilibrium) even more likely to occur. In practice, Amazon has become by far the most dominant retail platform with substantial monopoly power (because of its already-established two-sided network), and even its rival platform Buy.com has begun selling as a third-party seller on Amazon. Future research may explicitly study, both empirically and analytically, the competition between competing platforms such as Amazon with eBay or Sears.com, or Apple's App Store with Google's Android Market.

Appendix for Chapter 3

Proof of Lemma 1: Without threat of entry from Amazon, the seller's optimal prices and service levels should be the same across the two periods since Amazon's fee \bar{f} does not change across periods. The seller chooses his service level ($\bar{e}_t^{(i)} \geq 0$) and price ($\bar{p}_t^{(i)} \geq 0$) to maximize his profit for each period i :

$$\bar{\Pi}_t^{(i)} = (\gamma_t + \bar{e}_t^{(i)} - b\bar{p}_t^{(i)}) [\bar{p}_t^{(i)} - c - \bar{f} - k(\bar{e}_t^{(i)})^2].$$

The first order conditions (FOC) are

$$\frac{\partial \bar{\Pi}_t^{(i)}}{\partial \bar{p}_t^{(i)}} = (\gamma_t + \bar{e}_t^{(i)} - b\bar{p}_t^{(i)}) - b [\bar{p}_t^{(i)} - c - \bar{f} - k(\bar{e}_t^{(i)})^2] = 0 \quad (\text{A1})$$

$$\frac{\partial \bar{\Pi}_t^{(i)}}{\partial \bar{e}_t^{(i)}} = (\gamma_t + \bar{e}_t^{(i)} - b\bar{p}_t^{(i)}) \cdot (-2k\bar{e}_t^{(i)}) + [\bar{p}_t^{(i)} - c - \bar{f} - k(\bar{e}_t^{(i)})^2] = 0 \quad (\text{A2}).$$

From (A1), we obtain $\bar{p}_t^{(i)} = \frac{\gamma_t + \bar{e}_t^{(i)} + b(c + \bar{f} + k(\bar{e}_t^{(i)})^2)}{2b}$. Substituting this into (A2), we then solve for $\bar{e}_t^{(i)}$:

$\frac{\partial \bar{\Pi}_t^{(i)}}{\partial \bar{e}_t^{(i)}} = k \left(\bar{e}_t^{(i)} - \frac{1}{2bk} \right) \left\{ bk \left(\bar{e}_t^{(i)} \right)^2 - \bar{e}_t^{(i)} - [\gamma_t - b(c + \bar{f})] \right\} = 0$. Thus, the potential FOC solutions are:

$$\bar{e}_t^{*(i)} = \frac{1}{2bk} \quad (\text{A3})$$

$$\bar{e}_t^{*(i)} = \frac{1 \pm \sqrt{1 + 4bk[\gamma_t - b(c + \bar{f})]}}{2bk} \quad (\text{A4})$$

We eliminate (A4) since simple algebra shows it yields zero demand: $\bar{q}_t^{(i)}(\bar{f}) = \gamma_t + \bar{e}_t^{*(i)} - b\bar{p}_t^{*(i)} = 0$.

Alternatively, one can formally apply the second partial derivative test (using the Hessian matrix) to show

that (A3) is the local maximum and (A4) is a saddle point. It is easy to show the boundary of $\bar{e}_t^{(i)} = 0$ and

$\bar{p}_t^{(i)} = 0$ yields a lower profit than (A3). Hence, (A3) is the seller's (global) profit-maximizing service

level with the corresponding optimal price and profit respectively given by (A5) and (A6) below.

$$\bar{p}_t^{*(i)}(\bar{f}) = \frac{\gamma_t + b(c + \bar{f}) + \frac{3}{4bk}}{2b} \quad (\text{A5})$$

$$\bar{\Pi}_t^*(\bar{f}) = \bar{\Pi}_t^{*(1)} + \bar{\Pi}_t^{*(2)} = \frac{[\gamma_t - b(c + \bar{f}) + \frac{1}{4bk}]^2}{2b} \quad (\text{A6})$$

Lemma 1 immediately follows. \square

Equilibrium Outcome When Amazon Commits to “No Entry” (Section 3.4.2)

The optimal service levels, prices and overall profit for a seller of type t are given by (A3), (A5) and (A6), respectively. Amazon's expected total profit is given by $\bar{\Pi}_A(\bar{f}) = \sum_{i=1}^2 \left\{ \bar{f} \left[\theta \bar{q}_H^{(i)}(\bar{f}) + (1-\theta) \bar{q}_L^{(i)}(\bar{f}) \right] \right\}$, where $\bar{q}_t^{(i)}(\bar{f}) = \gamma_t + \bar{e}_t^{*(i)} - b \bar{p}_t^{*(i)}(\bar{f})$ is type $t \in \{L, H\}$ seller's quantity sold in period i as a function of \bar{f} . Substitution of (A3) and (A5) into $\bar{q}_t^{(i)}(\bar{f})$ leads to $\bar{\Pi}_A(\bar{f}) = \bar{f} \left[\theta \gamma_H + (1-\theta) \gamma_L - b(c + \bar{f}) + \frac{1}{4bk} \right]$. Amazon's equilibrium fee and profit are:²⁸

$$\bar{f}^* = \frac{\theta \gamma_H + (1-\theta) \gamma_L - bc + \frac{1}{4bk}}{2b} \quad (\text{A7})$$

$$\bar{\Pi}_A^* = \frac{\left[\theta \gamma_H + (1-\theta) \gamma_L - bc + \frac{1}{4bk} \right]^2}{4b} \quad (\text{A8})$$

Substituting (A7) into (A5) and (A6), we obtain the equilibrium outcome for type $t \in \{L, H\}$ seller.

$$\bar{p}_t^{*(1)} = \bar{p}_t^{*(2)} = \frac{2\gamma_t + \theta \gamma_H + (1-\theta) \gamma_L + bc + \frac{7}{4bk}}{4b} \quad (\text{A9})$$

$$\bar{\Pi}_t^* = \frac{\left[2\gamma_t - \theta \gamma_H - (1-\theta) \gamma_L - bc + \frac{1}{4bk} \right]^2}{8b} \quad (\text{A10})$$

Separating Outcome

Amazon's expected profit is expressed below:

$$\Pi_{A,sep}(f_{sep}, e_A, p_A) = \theta f_{sep} \cdot \frac{\gamma_H - b(c + f_{sep}) + \frac{1}{4bk}}{2} + \theta [(p_A - c - k e_A^2)(\gamma_H + e_A - b p_A) - F].$$

A proof very similar to that for Lemma 1 shows that $e_A^* = \frac{1}{2bk}$ and $p_A^* = \frac{\gamma_H + bc + \frac{3}{4bk}}{2b}$. With this, we

can rewrite Amazon's expected profit as a function of only f_{sep} where, as discussed before, $\frac{\gamma_L - bc + \frac{1}{4bk}}{b} \leq$

$$f_{sep} < \frac{\gamma_H - bc + \frac{1}{4bk}}{b}.$$

²⁸ Here we have implicitly assumed a non-boundary solution. That is, θ and γ_H are not both so large that Amazon will totally ignore the possibility of an L-type seller and target only the H-type seller (by charging $f^* = \frac{\gamma_H - bc + \frac{1}{4bk}}{2b}$).

One can easily show that such a boundary solution requires both $\theta > \left(\frac{\gamma_L - bc + \frac{1}{4bk}}{\gamma_H - \gamma_L} \right)^2$ and $\gamma_H > 2\gamma_L - bc + \frac{1}{4bk}$.

$$\Pi_{A,sep}(f_{sep}) = \theta f_{sep} \cdot \frac{\gamma_H - b(c + f_{sep}) + \frac{1}{4bk}}{2} + \theta \left[\frac{(\gamma_H - bc + \frac{1}{4bk})^2}{4b} - F \right] \quad (A11)$$

Recall that Assumption C1(ii) implies that if Amazon knows that the seller is an H-type, it will enter the market in the second period, since doing so yields a profit of $\frac{(\gamma_H - bc + \frac{1}{4bk})^2}{4b} - F$ rather than the maximum potential profit of $\frac{(\gamma_H - bc + \frac{1}{4bk})^2}{8b}$ from the fee collected from the H-type. At this point, we cannot yet fully determine whether or not a separating equilibrium will be realized for the overall game, because we must also calculate Amazon's profit for any fee $f < \frac{\gamma_L - bc + \frac{1}{4bk}}{b}$. After we fully specify Amazon's expected profit for all fee intervals, we then determine which fee maximizes Amazon's expected profit and hence which type of equilibrium is realized. Note that by maximizing (A11), *if* a separating equilibrium is realized, Amazon's fee must be given by (A12), its expected profit by (A13), and the seller's profit by (A14)-(A15). The two forms in (A12)-(A14) are according to whether the maximum occurs at an interior point or at the boundary. Later, we specify in Proposition 1 the condition under which this equilibrium is realized.

$$f_{sep}^* = \begin{cases} \frac{\gamma_L - bc + \frac{1}{4bk}}{b}, & \text{if } \gamma_H < 2\gamma_L - bc + \frac{1}{4bk}, \\ \frac{\gamma_H - bc + \frac{1}{4bk}}{2b}, & \text{if } \gamma_H \geq 2\gamma_L - bc + \frac{1}{4bk}, \end{cases} \quad (A12)$$

$$\Pi_{A,sep}^* = \begin{cases} \theta \left[\frac{(\gamma_L - bc + \frac{1}{4bk})(\gamma_H - \gamma_L)}{2b} + \frac{(\gamma_H - bc + \frac{1}{4bk})^2}{4b} - F \right], & \text{if } \gamma_H < 2\gamma_L - bc + \frac{1}{4bk}, \\ \theta \left[\frac{3(\gamma_H - bc + \frac{1}{4bk})^2}{8b} - F \right], & \text{if } \gamma_H \geq 2\gamma_L - bc + \frac{1}{4bk} \end{cases} \quad (A13)$$

$$\Pi_{H,sep}^* = \begin{cases} \frac{(\gamma_H - \gamma_L)^2}{4b}, & \text{if } \gamma_H < 2\gamma_L - bc + \frac{1}{4bk} \\ \frac{(\gamma_H - bc + \frac{1}{4bk})^2}{16b}, & \text{if } \gamma_H \geq 2\gamma_L - bc + \frac{1}{4bk} \end{cases} \quad (A14)$$

$$\Pi_{L,sep}^* = 0 \quad (A15)$$

Pooling Outcome

For the sake of clarity and completeness, we list below the seller's pooling equilibrium decisions (as indicated by the subscript "pool") for both periods with (A16) and (A17) for the L-type seller and (A18)-(A21) for the H-type seller.

$$e_{L,pool}^{*(1)} = e_{L,pool}^{*(2)} = \frac{1}{2bk} \quad (A16)$$

$$p_{L,pool}^{*(1)} = p_{L,pool}^{*(2)} = \frac{\gamma_L + b(c+f) + \frac{3}{4bk}}{2b} \quad (A17)$$

$$e_{H,pool}^{*(1)} = \gamma_L + \frac{1}{2bk} - \gamma_H \quad (A18)$$

$$p_{H,pool}^{*(1)} = \frac{\gamma_L + b(c+f) + \frac{3}{4bk}}{2b} \quad (A19)$$

$$e_{H,pool}^{*(2)} = \frac{1}{2bk} \quad (A20)$$

$$p_{H,pool}^{*(2)} = \frac{\gamma_H + b(c+f) + \frac{3}{4bk}}{2b} \quad (A21)$$

From the above, we can easily compute the overall profit for each type of seller as a function of f .

$$\Pi_{L,pool}^*(f) = \frac{\left[\gamma_L - b(c+f) + \frac{1}{4bk}\right]^2}{2b} \quad (A22)$$

$$\begin{aligned} \Pi_{H,pool}^*(f) &= \Pi_{H,pool}^{*(1)}(f) + \Pi_{H,pool}^{*(2)}(f) \\ &= \left[p_{H,pool}^{*(1)} - c - f - k(e_{H,pool}^{*(1)})^2\right] \left(\gamma_H + e_{H,pool}^{*(1)} - bp_{H,pool}^{*(1)}\right) + \frac{\left[\gamma_H - b(c+f) + \frac{1}{4bk}\right]^2}{4b} \end{aligned} \quad (A23)$$

Amazon's expected total profit is given by $\Pi_{A,pool}(f) = \Pi_{A,pool}^{(1)}(f) + \Pi_{A,pool}^{(2)}(f)$. It chooses f to maximize its expected profit. Amazon's optimal pooling fee and profit are given by:²⁹

$$f_{pool}^* = \frac{\frac{\theta(\gamma_H - \gamma_L)}{2} + \gamma_L - bc + \frac{1}{4bk}}{2b} \quad (A24)$$

$$\Pi_{A,pool}^* = \frac{\left[\frac{\theta(\gamma_H - \gamma_L)}{2} + \gamma_L - bc + \frac{1}{4bk}\right]^2}{4b} \quad (A25)$$

²⁹ In the proof of Proposition 1, we show that Amazon's pooling profit at the boundary $f \rightarrow \frac{\gamma_L - bc + \frac{1}{4bk}}{b}$ is always lower than its separating profit. Hence, (A25) is the only possible pooling equilibrium fee.

Proof of Proposition 1: Note that Amazon dictates which equilibrium is realized by selecting the appropriate fee f . Thus, for any given set of parameter values, we simply need to compare Amazon's expected profits under the two types of equilibria (in two different fee intervals) to determine the realized equilibrium for the overall game. For any given set of parameter values, we obtain a separating equilibrium if $\Pi_{A,sep}^* \geq \Pi_{A,pool}^*$ and a pooling equilibrium if $\Pi_{A,sep}^* < \Pi_{A,pool}^*$.

Case 1: $\gamma_H \geq 2\gamma_L - bc + \frac{1}{4bk}$

We examine two sub-cases. First, we consider the condition of $\theta > \frac{2(\gamma_L - bc + \frac{1}{4bk})}{\gamma_H - \gamma_L}$. As discussed before, a pooling equilibrium requires $f < \frac{\gamma_L - bc + \frac{1}{4bk}}{b}$ such that the L-type will profitably sell a positive quantity. But if $\theta > \frac{2(\gamma_L - bc + \frac{1}{4bk})}{\gamma_H - \gamma_L}$, the FOC solution $f_{pool}^* = \frac{\frac{\theta(\gamma_H - \gamma_L)}{2} + \gamma_L - bc + \frac{1}{4bk}}{2b} > \frac{\gamma_L - bc + \frac{1}{4bk}}{b}$ implies that Amazon's best pooling outcome corresponds to a fee $f \rightarrow \frac{\gamma_L - bc + \frac{1}{4bk}}{b}$. As $\rightarrow \frac{\gamma_L - bc + \frac{1}{4bk}}{b}$, $\Pi_{A,pool}(f) \rightarrow \frac{\theta(\gamma_H - \gamma_L)[\gamma_L - bc + \frac{1}{4bk}]}{2b}$. From (A13) and Assumption C1(ii), we get $\Pi_{A,sep}^* = \theta \left[\frac{3(\gamma_H - bc + \frac{1}{4bk})^2}{8b} - F \right] > \frac{\theta(\gamma_H - bc + \frac{1}{4bk})^2}{4b} > \frac{\theta(\gamma_H - \gamma_L)[\gamma_L - bc + \frac{1}{4bk}]}{2b}$. Note that the last inequality is proved by expanding the terms:

$$\frac{\theta(\gamma_H - bc + \frac{1}{4bk})^2}{4b} > \frac{\theta(\gamma_H - \gamma_L)[\gamma_L - bc + \frac{1}{4bk}]}{2b} \Leftrightarrow (\gamma_H - \gamma_L)^2 - \left(\gamma_L - bc + \frac{1}{4bk} \right)^2 > 0 \Leftrightarrow \gamma_H - \gamma_L > \gamma_L - bc + \frac{1}{4bk},$$

which is true under Case 1. Thus, we conclude the separating equilibrium is realized if $\theta > \frac{2(\gamma_L - bc + \frac{1}{4bk})}{\gamma_H - \gamma_L}$.

Second, we consider the condition of $\theta \leq \frac{2(\gamma_L - bc + \frac{1}{4bk})}{\gamma_H - \gamma_L}$, which means the best pooling outcome occurs at the FOC point (A24). We compare Amazon's profits from the potential separating equilibrium (A13) and the potential pooling equilibrium (A25). The separating equilibrium is realized if and only if

$$\Pi_{A,sep}^* \geq \Pi_{A,pool}^* \text{ or } \theta \left[\frac{3(\gamma_H - bc + \frac{1}{4bk})^2}{8b} - F \right] - \frac{\left[\frac{\theta(\gamma_H - \gamma_L)}{2} + \gamma_L - bc + \frac{1}{4bk} \right]^2}{4b} \geq 0.$$




Figure A-1: Amazon's pooling and separating profits

We plot both $\Pi_{A,sep}^*$ and $\Pi_{A,pool}^*$ as a function of θ in Figure A-1. These two curves (one linear and one quadratic in θ) have two points of intersection with one intersection on each side of $\theta = 1$, because

$$\begin{aligned}
 \lim_{\theta \rightarrow 1} \{ \Pi_{A,sep}^* - \Pi_{A,pool}^* \} &= \left[\frac{3(\gamma_H - bc + \frac{1}{4bk})^2}{8b} - F \right] - \frac{\left[\frac{(\gamma_H - \gamma_L)}{2} + \gamma_L - bc + \frac{1}{4bk} \right]^2}{4b} \\
 &= \frac{1}{8b} \left[3 \left(\gamma_H - bc + \frac{1}{4bk} \right)^2 - 2 \left(\frac{\gamma_H + \gamma_L}{2} - bc + \frac{1}{4bk} \right)^2 - 8bF \right] \\
 &> \frac{1}{8b} \left[3 \left(\gamma_H - bc + \frac{1}{4bk} \right)^2 - 2 \left(\frac{\gamma_H + \gamma_L}{2} - bc + \frac{1}{4bk} \right)^2 - 8b \frac{(\gamma_H - bc + \frac{1}{4bk})^2}{8b} \right], \text{ from assumption C1(ii).} \\
 &= \frac{1}{4b} \left[\left(\gamma_H - bc + \frac{1}{4bk} \right)^2 - \left(\frac{\gamma_H + \gamma_L}{2} - bc + \frac{1}{4bk} \right)^2 \right] > 0.
 \end{aligned}$$

Thus, if $\theta \geq \theta^*$, the separating equilibrium will be realized; if $\theta < \theta^*$, the pooling equilibrium is realized,

where θ^* , the smaller root of $\theta \left[\frac{3(\gamma_H - bc + \frac{1}{4bk})^2}{8b} - F \right] = \frac{\left[\frac{\theta(\gamma_H - \gamma_L)}{2} + \gamma_L - bc + \frac{1}{4bk} \right]^2}{4b}$, is given by $\theta^* =$

$$\frac{3(\gamma_H - bc + \frac{1}{4bk})^2 - 8bF - 2(\gamma_H - \gamma_L)(\gamma_L - bc + \frac{1}{4bk}) - \sqrt{\left[3(\gamma_H - bc + \frac{1}{4bk})^2 - 8bF - 2(\gamma_H - \gamma_L)(\gamma_L - bc + \frac{1}{4bk}) \right]^2 - 4(\gamma_H - \gamma_L)^2(\gamma_H - bc + \frac{1}{4bk})^2}}{(\gamma_H - \gamma_L)^2}$$

Case 2: $\gamma_H < 2\gamma_L - bc + \frac{1}{4bk}$

The separating equilibrium is realized if and only if $\Pi_{A,sep}^* \geq \Pi_{A,pool}^*$ or $\theta \left[\frac{(\gamma_L + \frac{1}{4bk} - bc)(\gamma_H - \gamma_L)}{2b} + \frac{(\gamma_H - bc + \frac{1}{4bk})^2}{4b} - F \right] - \frac{[\frac{\theta(\gamma_H - \gamma_L)}{2} + \gamma_L - bc + \frac{1}{4bk}]^2}{4b} \geq 0$. Thus, if $\theta \geq \theta^*$, the separating equilibrium will be realized; if $\theta < \theta^*$, the pooling equilibrium is realized, where $\theta^* =$

$$\frac{2(\gamma_H - bc + \frac{1}{4bk})^2 - 8bF + 2(\gamma_H - \gamma_L)(\gamma_L - bc + \frac{1}{4bk}) - \sqrt{[2(\gamma_H - bc + \frac{1}{4bk})^2 - 8bF + 2(\gamma_H - \gamma_L)(\gamma_L - bc + \frac{1}{4bk})]^2 - 4(\gamma_H - \gamma_L)^2(\gamma_H - bc + \frac{1}{4bk})^2}}{(\gamma_H - \gamma_L)^2}$$

Proof of Proposition 2: We define θ^* to be the cutoff constant given in Proposition 1. Hence, if Amazon keeps its entry option and $\theta < \theta^*$, the pooling equilibrium outcome will be realized. Comparing (A7) and (A24), it is simple to show $f_{pool}^* < \bar{f}^*$. Comparing Amazon's pooling equilibrium profit given by (A25) with its no-threat-of-entry profit given by (A8), one easily shows that $\Pi_{A,pool}^* < \bar{\Pi}_A^*$. \square

Proof of Proposition 3: Please refer to Technical Appendix for proof.

Proof of Proposition 4: Computing consumer surplus (CS) from an inverse linear demand function is straightforward. The second-period consumer surplus depends on the seller's type. From equations (A3) and (A5) and the demand function, one easily shows that without Amazon's threat of entry, the consumer surplus in each period for the t -type seller is given by $\overline{CS}_t^{(1)}(\bar{f}^*) = \overline{CS}_t^{(2)}(\bar{f}^*) = \frac{[\gamma_t - b(c + \bar{f}^*) + \frac{1}{4bk}]^2}{8b}$. When Amazon keeps its entry option, the pooling equilibrium is realized if $\theta < \theta^*$; the second-period consumer surplus is $CS_t^{(2)}(f_{pool}^*) = \frac{[\gamma_t - b(c + f_{pool}^*) + \frac{1}{4bk}]^2}{8b}$. Since $f_{pool}^* < \bar{f}^*$, we conclude $CS_t^{(2)}(f_{pool}^*) > \overline{CS}_t^{(2)}(\bar{f}^*)$.

For the first period, $CS_L^{(1)}(f_{pool}^*) > \overline{CS}_L^{(1)}(\bar{f}^*)$ obviously holds since in the first period, the L-type seller chooses the same first-best price and service level as in the second period. Under the threat of entry, the pooling consumer surplus in the first period is exactly the same for both types of sellers, since the H-type exactly mimics the L-type. So, $CS_H^{(1)}(f_{pool}^*) = CS_L^{(1)}(f_{pool}^*) = \frac{[\gamma_L - b(c + f_{pool}^*) + \frac{1}{4bk}]^2}{8b}$. Using equations (A7) and (A24), one can then easily show $CS_H^{(1)}(f_{pool}^*) < \overline{CS}_H^{(1)}(\bar{f}^*)$. \square

Proof of Proposition 5: This follows immediately from (A24) and (A12).

Proof of Proposition 6:

$$f_{rev}^* - f_{pool}^* = \frac{\theta(\gamma_H - bc + \frac{1}{4bk}) + 2(1-\theta)(\gamma_L - bc + \frac{1}{4bk})}{2b(2-\theta)} - \frac{\frac{\theta(\gamma_H - \gamma_L)}{2} + \gamma_L - bc + \frac{1}{4bk}}{2b} = \frac{\theta^2(\gamma_H - \gamma_L)}{4b(2-\theta)} > 0.$$

$$\begin{aligned} \text{If } \gamma_H < 2\gamma_L - bc - \frac{1}{4bk}, f_{rev}^* - f_{sep}^* &= \frac{\theta(\gamma_H - bc + \frac{1}{4bk}) + 2(1-\theta)(\gamma_L - bc + \frac{1}{4bk})}{2b(2-\theta)} - \frac{\gamma_L - bc + \frac{1}{4bk}}{b} \\ &= \frac{\theta(\gamma_H - bc + \frac{1}{4bk}) + 2(1-\theta)(\gamma_L - bc + \frac{1}{4bk}) - 2(2-\theta)(\gamma_L - bc + \frac{1}{4bk})}{2b(2-\theta)} = \frac{\theta(\gamma_H - bc + \frac{1}{4bk}) - 2(\gamma_L - bc + \frac{1}{4bk})}{2b(2-\theta)} < 0. \end{aligned}$$

$$\begin{aligned} \text{If } \gamma_H \geq 2\gamma_L - bc - \frac{1}{4bk}, f_{rev}^* - f_{sep}^* &= \frac{\theta(\gamma_H - bc + \frac{1}{4bk}) + 2(1-\theta)(\gamma_L - bc + \frac{1}{4bk})}{2b(2-\theta)} - \frac{\gamma_H - bc + \frac{1}{4bk}}{2b} \\ &= \frac{\theta(\gamma_H - bc + \frac{1}{4bk}) + 2(1-\theta)(\gamma_L - bc + \frac{1}{4bk}) - (2-\theta)(\gamma_H - bc + \frac{1}{4bk})}{2b(2-\theta)} = \frac{(1-\theta)(\gamma_L - \gamma_H)}{b(2-\theta)} < 0. \end{aligned}$$

Here we have implicitly assumed the interesting case of $\theta < \frac{2(\gamma_L - bc + \frac{1}{4bk})}{\gamma_H - bc + \frac{1}{4bk}}$ such that both types of sellers are targeted by Amazon. □

Technical Appendix for Chapter 3

Equilibrium Refinement

To refine the infinite continuum of pooling equilibria supported by different out-of-equilibrium beliefs, we proceed in two main steps. First, we use the intuitive criterion to rule out certain pooling outcomes. However, there remains a range of pooling outcomes that cannot be refined further using the intuitive criterion or even the stronger universal divinity criterion (Banks and Sobel 1987). Therefore, in the second step, we apply additional logical arguments that are applicable in our game and help us to pin down the unique pooling equilibrium.

We now proceed with the first step of applying the intuitive criterion. Let Γ^* denote the overall demand intercept that Amazon observes at the pooling equilibrium under consideration. Consider a pooling equilibrium in which, conditional on f and Γ^* , both types of sellers pool at a price $p^\#$ different from the L-type's first-best price. Note that $p^\#$ cannot be smaller than the L-type's first-best price, because charging such a low price is a dominated strategy for *both* types of sellers. Furthermore, if $p^\#$ is larger than the L-type's first-best price, then the H-type seller is worse off by deviating to the L-type's first-best price while the L-type is better off with such a deviation. Thus, the equilibrium fails the intuitive criterion because the seller's deviation to the L-type's first-best price tells Amazon that the deviating seller must be an L-type since an H-type will not voluntarily choose to make such an unprofitable deviation. Therefore we conclude that for any pooling equilibrium to survive the intuitive criterion, both types of sellers must choose the L-type seller's first-best price conditional on f and Γ^* .

Now we consider a pooling equilibrium with $\Gamma^* > \gamma_L + \frac{1}{2bk}$. This implies that $e_L^{*(1)} > \frac{1}{2bk}$, i.e., both types of sellers are pooling at a service level larger than the L-type's first-best service. Clearly, the L-type seller makes a higher profit by deviating to his first-best service level of $\frac{1}{2bk}$ creating a lower overall demand intercept of $\Gamma = \gamma_L + \frac{1}{2bk}$. If such a deviation also makes the H-type seller better off, then the original outcome cannot be at equilibrium because Γ^* corresponds to a dominated strategy for both types of sellers. Furthermore, if such a deviation makes the H-type seller worse off, then the pooling

equilibrium fails the intuitive criterion—the fact that the seller is voluntarily reducing Γ serves to inform Amazon that he must be an L-type and hence breaks the equilibrium. Thus, no pooling equilibrium with $\Gamma^* > \gamma_L + \frac{1}{2bk}$ survives the intuitive criterion. Further, since the demand intercept for the H-type seller is at least γ_H , there cannot be any pooling equilibrium with $\Gamma^* < \gamma_H$. Therefore, we must have $\gamma_H \leq \Gamma^* \leq \gamma_L + \frac{1}{2bk}$.

For each Γ^* satisfying $\gamma_H \leq \Gamma^* \leq \gamma_L + \frac{1}{2bk}$, however, the belief $\mu(H|\Gamma) = \begin{cases} 1, & \text{if } \Gamma > \Gamma^*, \\ \theta, & \text{if } \gamma_H \leq \Gamma \leq \Gamma^*, \\ 0, & \text{if } \Gamma < \gamma_H \end{cases}$

supports a pooling equilibrium outcome that survives the intuitive criterion as well as the stronger universal divinity criterion. The intuitive reason is that for the above criteria to eliminate an equilibrium, at least one type of seller must have a profitable deviation from the equilibrium signal. But for an equilibrium with $\gamma_H \leq \Gamma^* \leq \gamma_L + \frac{1}{2bk}$ and given the belief above, *both* types of sellers will make a strictly lower profit with any out-of-equilibrium signal. Deviating to any $\Gamma < \Gamma^*$ will not trigger Amazon's entry but yields lower profits for both types of sellers. Deviating to any $\Gamma > \Gamma^*$ will trigger Amazon's entry, also leading to a lower profit for the seller. This is because the seller's maximum profit by sending an out-of-equilibrium signal $\Gamma > \Gamma^*$ (which comes only from the first period since such a signal triggers Amazon's entry) is the same as the seller's second-period profit alone if Amazon does not enter. So, the question remains which of these remaining pooling equilibria (Γ^*) is the most “reasonable” equilibrium outcome?

We now proceed with the second step to show that $\Gamma^* = \gamma_L + \frac{1}{2bk}$ is the only reasonable equilibrium. It is straightforward to show that given $\gamma_H \leq \Gamma^* \leq \gamma_L + \frac{1}{2bk}$, the pooling price must be

$p_{t,pool}^{(1)}(\Gamma^*) = \frac{\Gamma^* + bk(\Gamma^* - \gamma_L)^2 + b(c+f)}{2b}$ and that the L-type and the H-type seller's pooling profits are given by

$$\Pi_{L,pool}(\Gamma^*) = \frac{[\Gamma^* - bk(\Gamma^* - \gamma_L)^2 - b(c+f)]^2}{4b} + \frac{[\gamma_L - b(c+f) + \frac{1}{4bk}]^2}{4b} \text{ and}$$

$$\Pi_{H,pool}(\Gamma^*) = \left[\frac{\Gamma^* + bk(\Gamma^* - \gamma_L)^2 - 2bk(\Gamma^* - \gamma_H)^2 - b(c+f)}{2b} \right] \left[\frac{\Gamma^* - bk(\Gamma^* - \gamma_L)^2 - b(c+f)}{2} \right] + \frac{[\gamma_H - b(c+f) + \frac{1}{4bk}]^2}{4b}, \text{ respectively.}$$

Further, Amazon's total profit from the pooling outcome is given by

$$\Pi_{A,pool}(\Gamma^*) = f \frac{\Gamma^* - bk(\Gamma^* - \gamma_L)^2 - b(c+f)}{2} + f \left[\theta \frac{\gamma_L - b(c+f) + \frac{1}{4bk}}{2} + (1 - \theta) \frac{\gamma_H - b(c+f) + \frac{1}{4bk}}{2} \right], \text{ where the first term}$$

is Amazon's first-period profit and the second term its second-period profit. One can also easily show that

$$\text{given } \Gamma^*, \text{ Amazon maximizes its profit with a fee of } f_{pool}^*(\Gamma^*) = \frac{\Gamma^* - bk(\Gamma^* - \gamma_L)^2 + \theta\gamma_H + (1-\theta)\gamma_L - 2bc + \frac{1}{4bk}}{4b}.$$

Note that for any $f > 0$, $\Pi_{L,pool}(\Gamma^*)$, $\Pi_{H,pool}(\Gamma^*)$, and $\Pi_{A,pool}(\Gamma^*)$ are *all* monotonically increasing functions on $\gamma_H \leq \Gamma^* \leq \gamma_L + \frac{1}{2bk}$. This suggests that $\Gamma^* = \gamma_L + \frac{1}{2bk}$ corresponds to the only reasonable belief by Amazon. That is, any pooling equilibrium with Amazon's belief corresponding to $\Gamma^* < \gamma_L + \frac{1}{2bk}$ is not reasonable. To see this, let us consider a deviation by Amazon from a pooling equilibrium with $\Gamma^* < \gamma_L + \frac{1}{2bk}$. Amazon's deviation from equilibrium suggests to the seller that it must believe that the seller's response to that deviation can make Amazon better off. The possible responses by both types of sellers that can make Amazon's deviation profitable correspond to some $\Gamma \in (\Gamma^*, \gamma_L + \frac{1}{2bk}]$. Note also that on this interval, the seller's response corresponding to $\Gamma = \gamma_L + \frac{1}{2bk}$ is not only the most profitable response for Amazon and but also for both types of sellers. Hence, Amazon will choose to deviate from the supposed equilibrium by setting a fee of $f_{pool}^*(\gamma_L + \frac{1}{2bk})$ instead of the equilibrium value of $f_{pool}^*(\Gamma^*)$. By doing so, Amazon convinces the seller that it anticipates his response will make Amazon better off. The seller (of both types) will choose, among all his responses that can make Amazon better off, the one corresponding to $\Gamma = \gamma_L + \frac{1}{2bk}$, which still yields a pooling outcome but gives the seller a larger profit and, in fact, the highest profit from all anticipated responses. This, in turn, justifies Amazon's belief about the seller's response to its deviation. That is, both parties will, with consistent beliefs about each other's rational responses, become better off with Amazon's deviation. Thus, Amazon will clearly choose

to deviate from the supposed equilibrium that corresponds to $\Gamma^* < \gamma_L + \frac{1}{2bk}$. We thereby conclude that the only reasonable (stable) pooling equilibrium must correspond to $\Gamma^* = \gamma_L + \frac{1}{2bk}$. \square

Proof of Proposition 3:

We define θ^* to be the cutoff constant given in Proposition 1. Hence, if Amazon keeps its entry option and $\theta < \theta^*$, the pooling equilibrium outcome will be realized. Substituting the pooling equilibrium fee (A24) into (A22) and (A23), we obtain the seller's profits.

$$\Pi_{L,pool}^* = \frac{\left[\gamma_L - b \left(c + \frac{\frac{\theta(\gamma_H - \gamma_L)}{2} + \gamma_L - bc + \frac{1}{4bk} \right) + \frac{1}{4bk} \right]^2}{2b} = \frac{\left[\gamma_L - bc + \frac{1}{4bk} - \frac{\theta(\gamma_H - \gamma_L)}{2} \right]^2}{8b} \quad (\text{A26})$$

$$\begin{aligned} \Pi_{H,pool}^* &= \frac{1}{4b} \left\{ \left[\gamma_L - b(c + f) + \frac{3}{4bk} - 2bk \left(\gamma_L - \gamma_H + \frac{1}{2bk} \right)^2 \right] \left(\gamma_L - b(c + f) + \frac{1}{4bk} \right) \right. \\ &\quad \left. + \left[\gamma_H - b(c + f) + \frac{1}{4bk} \right]^2 \right\} \\ &= \frac{1}{4b} \left\{ \left[\gamma_L - b \left(c + \frac{\frac{\theta(\gamma_H - \gamma_L)}{2} + \gamma_L - bc + \frac{1}{4bk} \right) + \frac{3}{4bk} - 2bk \left(\gamma_L - \gamma_H + \frac{1}{2bk} \right)^2 \right] \left(\gamma_L - \right. \right. \\ &\quad \left. \left. b \left(c + \frac{\frac{\theta(\gamma_H - \gamma_L)}{2} + \gamma_L - bc + \frac{1}{4bk} \right) + \frac{1}{4bk} \right) + \left[\gamma_H - b \left(c + \frac{\frac{\theta(\gamma_H - \gamma_L)}{2} + \gamma_L - bc + \frac{1}{4bk} \right) + \frac{1}{4bk} \right]^2 \right\} \\ &= \frac{1}{4b} \left\{ \left[\frac{\gamma_L - bc + \frac{1}{4bk}}{2} - \frac{\theta(\gamma_H - \gamma_L)}{4} + \frac{1}{2bk} - 2bk \left(\gamma_L - \gamma_H + \frac{1}{2bk} \right)^2 \right] \left(\frac{\gamma_L - bc + \frac{1}{4bk}}{2} - \frac{\theta(\gamma_H - \gamma_L)}{4} \right) + \left[\gamma_H - \frac{\gamma_L}{2} - \right. \right. \\ &\quad \left. \left. \frac{\theta(\gamma_H - \gamma_L)}{4} - \frac{bc}{2} + \frac{1}{8bk} \right]^2 \right\} \\ &= \frac{1}{16b} \left\{ \left[\gamma_L - bc + \frac{5}{4bk} - \frac{\theta(\gamma_H - \gamma_L)}{2} - 4bk \left(\gamma_L - \gamma_H + \frac{1}{2bk} \right)^2 \right] \left(\gamma_L - bc + \frac{1}{4bk} - \frac{\theta(\gamma_H - \gamma_L)}{2} \right) + \right. \\ &\quad \left. \left[(\gamma_H - \gamma_L) \left(2 - \frac{\theta}{2} \right) + \gamma_L - bc + \frac{1}{4bk} \right]^2 \right\} \quad (\text{A27}) \end{aligned}$$

Without threat of entry by Amazon, the seller's equilibrium profit is given by (A10). More specifically,

$$\bar{\Pi}_L^* = \frac{\left[\gamma_L - bc + \frac{1}{4bk} - \theta(\gamma_H - \gamma_L) \right]^2}{8b} \quad (\text{A28})$$

$$\bar{\Pi}_H^* = \frac{[(2-\theta)(\gamma_H - \gamma_L) + \gamma_L - bc + \frac{1}{4bk}]^2}{8b} \quad (\text{A29}).$$

Note that as seen in the proof of Proposition 1, $\gamma_L - bc + \frac{1}{4bk} - \theta(\gamma_H - \gamma_L) \geq 0$ must be true if the pooling equilibrium is realized. Thus, from (A26) and (A28), we obtain $\Pi_{L,pool}^* > \bar{\Pi}_L^*$. This result is intuitively sound since in both scenarios, the L-type seller picks the same service level and sets his first-best price given Amazon's fee but $f_{pool}^* < \bar{f}^*$.

From (A27) and (A29), we define a cubic polynomial function:

$$\begin{aligned} Z(\gamma_H) &\equiv 16b(\Pi_{H,pool}^* - \bar{\Pi}_H^*) \\ &= \left[\gamma_L - bc + \frac{5}{4bk} - \frac{\theta(\gamma_H - \gamma_L)}{2} - 4bk \left(\gamma_L - \gamma_H + \frac{1}{2bk} \right)^2 \right] \left(\gamma_L - bc + \frac{1}{4bk} - \frac{\theta(\gamma_H - \gamma_L)}{2} \right) + \left[(\gamma_H - \gamma_L) \left(2 - \frac{\theta}{2} \right) + \gamma_L - bc + \frac{1}{4bk} \right]^2 - 2 \left[(2 - \theta)(\gamma_H - \gamma_L) + \gamma_L - bc + \frac{1}{4bk} \right]^2 \end{aligned}$$




Figure A-2: Three solutions to $Z(\gamma_H) = 0$.

So, if $Z(\gamma_H) > 0$, the H-type seller is better off at the pooling equilibrium under Amazon's threat of entry than without Amazon's threat of entry. Figure A-2 illustrates the three intersections of this cubic polynomial function with the γ_H -axis (i.e., the three solutions to $Z(\gamma_H) = 0$). Note that $Z(\gamma_L) = 0$, so we can factor out $(\gamma_H - \gamma_L)$ from $Z(\gamma_H)$ and easily solve the other two solutions. It is straightforward albeit somewhat tedious to show that the smallest intersection is at $\gamma_H = \gamma_L$ and the largest is at some $\gamma_H > \gamma_L + \frac{1}{2bk}$ by showing $Z'(\gamma_L) > 0$ and $Z\left(\gamma_L + \frac{1}{2bk}\right) = -\left[\frac{(1-\theta)(\gamma_L - bc)}{bk} + \frac{3\theta^2 - 12\theta + 10}{8b^2k^2}\right] < 0$. Recall also that the meaningful region for γ_H is specified by Assumption C2. Therefore, Proposition 4 immediately follows where γ^* is the middle solution given below.

$$\gamma^* = \frac{1}{8bk\theta} \left\{ 10 - 8b^2ck - 8\theta + 3\theta^2 + 8bk\gamma_L(1 + \theta) - \left[100 + 64b^4c^2k^2 - 128b^3ck^2\gamma_L - 160\theta + 108\theta^2 - 48\theta^3 + 9\theta^4 - 16bk\gamma_L(\theta^2 + 8\theta - 10) + 16b^2k \left(4k\gamma_L^2 + c(\theta^2 + 8\theta - 10) \right) \right]^{\frac{1}{2}} \right\} \quad \square$$

Pooling outcome under nonlinear fee structure

We show the pooling outcome survives even if we adopt a nonlinear fee in our model. Suppose that Amazon adopts a nonlinear contract. Let the function $x(q)$ be Amazon's non-linear fee (based on q , the seller's quantity sold). If the non-linear fee, however it may be structured, is set such that the low-type seller does not find it profitable to enter the market (because it will sell zero units even if it enters), then we obtain a separating outcome—the high-type seller will offer its first-best service in the first period (conditional on the fee structure) and Amazon will replace the high-type seller in the second period.

If the nonlinear fee is such that the low-type seller sells some positive quantity, then we will obtain a pooling outcome provided that it is possible for the high-type seller to mask as a low-type, i.e., $\gamma_H \leq \gamma_L + e_L^*$, where e_L^* is the low-type seller's optimal service level conditional on Amazon's fee structure $x(q)$. Note that if the high-type seller does not mask his demand in the first period and gets identified by Amazon, he will make some profit $\Pi_{H,sep}^{(1)}$ in the first period but zero profit in the second period (since Amazon will enter the market). If he masks himself as a low-type, however, Amazon will not enter the market in the second period and the high-type seller's second-period profit will be at least as large as $\Pi_{H,sep}^{(1)}$. Since he also makes some profit in the first period, the high-type seller strictly prefers mimicking the low-type seller in the first period (by providing lower service), to prevent Amazon from learning his true type hence voiding Amazon's entry in the second period. Therefore, if Amazon's non-linear fee, however it may be structured, allows for the low-type seller's profitable entry, then the high-type seller prefers the pooling outcome to the separating outcome.

Now we show that, under some conditions, Amazon will also prefer the pooling outcome to the separating outcome even when it can choose any non-linear contract $x(q)$. Note that Amazon's separating

profit under any nonlinear contract is at most the high-type seller's maximum profit for the case of *zero* fee and *without* Amazon's threat of entry. That is, Amazon's maximum expected profit from a separating outcome must satisfy

$$\max [\Pi_{A,sep}|_{x(q)}] \leq \theta \cdot \max [\bar{\Pi}_H|_{x(q)=0}] = \theta \frac{[\gamma_H - bc + \frac{1}{4bk}]^2}{2b}, \text{ where the expression for the upper bound is obtained from equation (A6) by setting } \bar{f} = 0.$$

Note that Amazon's maximum pooling profit under a general (optimally chosen) nonlinear fee structure must be at least as large as its optimal pooling profit under the simple linear fee structure (which we studied in our original model). That is, Amazon's maximum expected profit from the pooling outcome under nonlinear contracts must satisfy

$$\max [\Pi_{A,pool}|_{x(q)}] \geq \max [\Pi_{A,pool}|_{x(q)=f_q}] = \frac{[\frac{\theta(\gamma_H - \gamma_L)}{2} + \gamma_L - bc + \frac{1}{4bk}]^2}{4b}, \text{ where the expression for the lower bound is given by equation (A25).}$$

Note that $\frac{[\frac{\theta(\gamma_H - \gamma_L)}{2} + \gamma_L - bc + \frac{1}{4bk}]^2}{4b} > \theta \frac{[\gamma_H - bc + \frac{1}{4bk}]^2}{2b}$ holds if $0 < \theta < \theta_c$, where θ_c is the smaller solution of

$$\text{the quadratic equation (in } \theta): \frac{[\frac{\theta(\gamma_H - \gamma_L)}{2} + \gamma_L - bc + \frac{1}{4bk}]^2}{4b} = \theta \frac{[\gamma_H - bc + \frac{1}{4bk}]^2}{2b}.$$

Therefore, $\max[\Pi_{A,pool}|_{x(q)}] > \max [\Pi_{A,sep}|_{x(q)}]$, i.e., Amazon also prefers the pooling outcome if $0 < \theta < \theta_c$.

Thus, even under nonlinear contracts, the pooling equilibrium still exists in some parameter region.

Below is Amazon's posterior belief that supports the pooling equilibrium:

$$\mu(H|p_t^{(1)}, q_t^{(1)}) = \begin{cases} 1, & \text{if } \Gamma > \gamma_L + e_L^*, \\ \theta, & \text{if } \gamma_H \leq \Gamma \leq \gamma_L + e_L^*, \text{ where the overall demand intercept } \Gamma \equiv q_t^{(1)} + bp_t^{(1)}. \\ 0, & \text{if } \Gamma < \gamma_H \end{cases} \square$$

4. Ad-Supported Platform in a Vertically Differentiated Market

4.1. Introduction

Software is a very significant industry in the modern economy; in 2006, the top 500 software firms earned a total revenue of \$380.3 billion.³⁰ Studying software firms' marketing and strategic decisions is clearly of practical importance. In recent years, many software firms have adopted the ad-supported software model. The ad-supported platform is an emerging, disruptive software licensing model in the traditional world of offline, perpetual software licensing; it allows software firms to earn revenues through online advertisements rather than or in addition to conventional software sales. Such software is typically free or sells at greatly reduced prices from its ad-free version. Some applications, e.g., most peer-to-peer software such as Kazaa, are offered only as free ad-supported software. For many firms such as Google, Microsoft, and many online game companies, the free versions of their applications display advertisements to users whereas their paid versions do not show any advertisements. Some firms actually derive most of their revenues from advertising rather than from sales of their software or services. The growing trend for ad-supported software is built upon some favorable characteristics of online advertising as compared to traditional advertising in the mass media. For example, online advertising is a real-time, two-way communication channel, which allows consumers to link directly to vendors' websites to make purchases or to acquire additional product information. Online advertising can offer dynamic targeting of consumers based on contexts or viewer profiles, and allows for real-time performance evaluations and advertisement changes. Such advantages of online advertising, the ubiquity of Internet access, and the fact that a fast growing number of consumers are spending more time online have enlarged the potential payoff of offering ad-supported software.

The rising importance of online advertising and the emergence of ad-supported software are bringing in fundamental changes in the software industry. Firms such as Yahoo and Google have been ad-supported from very early on, whereas the software giant Microsoft has only recently embraced the ad-supported model offering ad-supported software to consumers and small businesses through its Windows

³⁰ <http://www.softwaremag.com/L.cfm?Doc=984-10/2006>

Live and Office Live. While presently most software applications are offered ad-free or offline, a growing number of firms are beginning to take advantage of the burgeoning online advertising market by adopting the software as a service (SaaS) and ad-supported platforms. Not only have many consumer applications been offered as ad-supported software or online services, but some ad-supported, business applications (e.g., Office Live) have also emerged and gained some acceptance especially among small firms. Salesforce.com—the leader of on-demand CRM applications—has announced that it will offer an ad-supported edition of its applications featuring Google AdWords.³¹ The extant research literature on this new phenomenon is lacking though the idea of ad-supported software has been abundantly mentioned in the popular business press. The mobile-computing platforms have shown tremendous growth. For example, in Apple Inc.'s App Store (for iPhone and iPad) and Google's Android Market (for competing AndroidOS-based smartphones and tablet computers), we see hundreds of thousands of free ad-supported applications as well as ad-free, paid applications, and also a firm often offers its application in both a paid ad-free version and a free ad-supported version. Even in hardware markets, some companies are beginning to test the ad-supported model. For example, Amazon.com Inc. has just announced that it will offer a cheaper ad-supported version of Kindle, its bestselling electronic reader. All aspects of the device are the same only that for the ad-supported model, special offers and sponsored screensavers will display on the Kindle's screensaver and at the bottom of the home screen (without interrupting reading on it).³² This may be the company's first key step in turning its Kindle eReader into an advertising platform. The ad-supported model is clearly an important phenomenon in these new technology and Internet-enabled markets. We provide an economic analysis of the ad-free and ad-supported platforms to shed light on the following research questions:

1. Should a firm (an application developer) offer its product ad-supported, ad-free, or both? What are the main driving factors for its decision?

³¹ <http://www.salesforce.com/company/news-press/press-releases/2007/06/070605.jsp>, June 5, 2007.

³² <http://www.usatoday.com/tech/products/2011-04-11-amazon-kindle-ads.htm>, April 11, 2011.

2. In a competitive, vertically differentiated market, do firms have less or more incentives to adopt the ad-supported platform? Which firm(s) will adopt which platform(s) under what conditions?

New software licensing models such as pay-per-use, open-source, SaaS and ad-supported platforms are disrupting the traditional world of fixed-fee, perpetual licensing. Our research on ad-supported software complements the extant literature on other software licensing models (e.g., Sundararajan 2004, Choudhary et al. 1998, Huang and Sundararajan 2005, Jiang et al. 2007, Economides and Katsamakas 2006, Lee and Mendelson 2008). Extant research on the ad-supported software platform is scant; there is some related research on media and broadcasting. Prasad et al. (2003) show, with a stylized model with two types of consumers, that it is generally optimal for a monopolist media provider to offer both subscription and advertisements. Riggins (2003) studies a monopolist's price and content quality choices for its fee-based website and the quality choice for its free ad-supported website. Fan et al (2007) studies how media content quality affects a monopolist's pricing and advertising levels. Dewan et al. (2003) find that it may be optimal for an ad-supported website to initially show fewer ads and more contents to attract more viewers in the future resulting in higher "discounted total traffic." Anderson and Coate (2005) show that, in the broadcasting industry, the equilibrium levels of advertising and amount of programming can be above or below socially optimal levels. Several studies (e.g., Gal-Or and Dukes 2003 and Peitz and Valletti 2008) have also found that ad-sponsored firms tend to offer less horizontally differentiated products. Casadesus-Masanell and Zhu (2009) study a vertically differentiated market with an ad-sponsored entrant and find that the incumbent is more likely to prefer competing through the subscription-based or the ad-sponsored model rather than both of these models. We also focus on a vertically differentiated market and study competitive firms' incentives in adopting the ad-supported platform in the software applications market. We find that in a vertically differentiated software market with intense competition, firms may have incentives to adopt the ad-supported platform even if their advertising revenue does not cover the fixed cost required for that platform.

The Internet-enabled ad-supported platform allows software firms to essentially change their traditional software business model to a “two-sided market” model. Firms sell their software to consumers in one market but can also sell ads to advertisers in their advertising market. There is a growing literature on two-sided markets (e.g. Rochet and Tirole 2003 and 2006, Parker and Van Alstyne 2005, Armstrong 2006). The two-sided market manifested in the current context differs from the above research. Most of this literature assumes the existence of some non-internalized (positive) externalities among end-users. In the context of the ad-supported software platform, however, there is no externality because the firm’s benefit from advertising via its software comes at the expense of annoyances to consumers, who fully internalize such disutility when making usage or purchase decisions. To the best of our knowledge, our research questions regarding ad-supported software have not been clearly addressed in the literature. We model a vertically differentiated market with heterogeneous consumers in terms of their valuation for software features and their distaste (or intolerance) for advertisements shown through the ad-supported software. We examine both monopoly and competitive markets in which both ad-free and ad-supported software may be present.

Our analysis of the monopoly market shows that, ignoring any fixed cost, it is generally sub-optimal for a monopolist to offer only ad-free software. Though a monopolist finds versioning of quality differentiated (ad-free) software to be sub-optimal (Jones and Mendelson 1998, Bhargava and Choudhary 2001), we show that it is optimal to offer both ad-free and ad-supported versions if some consumers may have high distaste for advertisements or if the per-user advertising rate is not very high. When the per-user advertising rate is high, the monopolist will find it optimal to offer only ad-supported software, either at a reduced price or for free.

In a duopoly market in which each firm can adopt only one platform, we find that, first, unless one firm’s product is far inferior to the other’s, both firms will be better off if either firm offers an ad-supported product than if neither does. Second, the high quality firm plays a dominant role in choosing the favorable platform. If the per-user advertising rate is above some threshold, the high quality firm will choose to offer free ad-supported software, forcing the low quality firm to abandon that platform to offer

ad-free software. But if the advertising rate is below that threshold, the high quality firm will offer ad-free software whereas the low quality firm may decide to adopt either platform depending on how high the per-user advertising rate is.

When both firms can potentially adopt multiple platforms, we find that, at equilibrium, the low quality firm offers only ad-free software whereas the high quality firm will offer both ad-free and ad-supported software unless the advertising rate is very close to zero *and* the low quality firm's product is very inferior. More interestingly, we find that even if *neither* firm earns any positive advertising revenue or only one firm does, it is possible for *both* firms to benefit from the availability of the ad-supported platform. We find that the ad-supported offering has a moderating effect on price competition between firms' ad-free offerings. Our analysis suggests that, in a quality differentiated software market with intense price competition, firms may have incentives to adopt the ad-supported platform even if they do not expect to make enough advertising revenue to fully cover the fixed cost associated with the ad-supported platform.

The rest of this chapter is organized as follows. In section 4.2, we introduce our analytical framework and study a monopolist's optimal decisions of whether to offer its software ad-free, ad-supported, or both. Section 4.3 first examines two competitive cases of a vertically differentiated software market—one in which both firms sell only ad-free software, and one in which one firm offers free ad-supported software while the other offers ad-free software. We then consider firms' competitive single-platform and multi-platform entry decisions. We conclude in section 4.4 with discussions of possible future research.

4.2. Monopoly Market

We analyze a monopoly market by first comparing ad-supported with ad-free software and then examining whether the monopolist will find it optimal to offer both. We assume, as is typical for information goods such as software, that firm's fixed costs are sunk and that its marginal cost is constant and normalized to zero. Our analysis can be easily applied to consider product introduction decisions with

fixed costs. We simply need to consider the potential benefits from the product offering(s) with the required fixed cost. If the fixed cost is higher than the potential profit gain, the firm should not introduce the product(s).

Ad-supported software typically establishes an Internet connection and either periodically shows advertisements (e.g., text-based or banner ads) while the software is in use or displays advertisements only when the software is first started. Such software can reside on many physical platforms such as personal computers and handheld wireless devices. The software firm will charge advertisers a fee based on some agreed-upon metrics (e.g., per impression or per click). Generally, the more users the ad-supported software has, the more the firm will be able to charge its advertisers. We assume that the firm will make an advertising revenue of r per user of its ad-supported software. We will refer r as the per-user advertising rate. Effectively, one can consider r to be the per-user advertising price at which the software firm sells its exclusive advertising rights to a third-party ads agency, which manages the advertising done via the software directly incurring associated costs and charging fees from advertisers. Note also that r includes all future advertising revenues or fees discounted accordingly. Essentially, r also signifies the effectiveness of advertising through the ad-supported software. If the software allows for well targeted advertising based on users' preferences/profiles and the contexts of the software use, we expect r to be higher than if advertisements can only be "blindly" shown.

Note that our model applies to the mobile computing settings such as Apple's App Store and Google's Android Market, where the platform (iOS or Android OS) is developed and provided by one technology company (the platform owner) while other software firms or application developers pay the platform owner a percentage of revenue (from either the sales of their ad-free application or the advertising fee charged with the ad-supported application). In the case of both Apple's App Store and Google's Android Market, the percentage fee charged is 30% of the application seller's revenue on the platform.³³

³³ <https://market.android.com/support/bin/answer.py?answer=112622> and <http://developer.apple.com/programs/ios/distribute.html>

If the software is free of any advertisement, the consumer places a valuation of μ on the software; if advertisements are shown, the consumer will see the software's value as impaired by an amount of θ . Such impairment perceived by consumers may also be thought of as their distaste/intolerance for advertising via the software or the additional cost/inconvenience of having the Internet connection. Consumers may have different valuations for software features and may be bothered by advertisements to a very different extent. That is, consumers are heterogeneous in terms of μ and θ . We assume consumers are uniformly distributed in these dimensions. We normalize μ to $\mu \sim \text{uniform}[0,1]$, and assume θ to be distributed independent of μ : $\theta \sim \text{uniform}[0,\lambda]$, where $\lambda > 0$ and the total market size is normalized to 1. The independence assumption is reasonable since, for instance, a high valuation user may have either high or low tolerance levels for advertisements displayed in the ad-supported software.

Table 4.1: Key Notations

Symbol	Explanation
μ	Consumers' valuation for software features, $\mu \sim \text{uniform}[0,1]$
θ	Consumers' distaste for advertisements, $\theta \sim \text{uniform}[0,\lambda]$, where $\lambda > 0$.
λ	Upper bound for consumers' distaste for advertisements in software
P_i	Price of product "i"
i	$i = "0"$ for ad-free, $i = "a"$ for ad-supported, $i = "1"$ for low quality, $i = "2"$ for high quality
u_i	Consumer's net utility from product "i"
Π_i	Profit function of firm i or product "i"
q_i	Quality level of product "i"
r	Per-user advertising rate

All key notations used in this chapter are summarized in Table 4.1. Put concisely, a consumer of type (μ, θ) derives a net utility of $u_0(\mu, P_0) = \mu - P_0$ from the ad-free software, and

$u_a(\mu, \theta, P_a) = \mu - P_a - \theta$ from the ad-supported software, where P_0 and P_a are the prices of the ad-free and ad-supported software, respectively. We use the subscript “0” to indicate that the software is ad-free (i.e., zero ads), and the subscript “a” to indicate that the software is ad-supported. We assume firms cannot set negative prices. Each consumer will use at most one type of software and derives zero utility if s/he does not use any. The profit maximizing monopolist decides whether to offer either ad-free or ad-supported software or both, and its corresponding decision variables are P_0 and P_a .

4.2.1. Ad-supported versus ad-free software

Now we compare the ad-free and ad-supported software platforms. We will later analyze the possibility of concurrently offering both. The monopolist can either sell an ad-free product at price P_0 or offer an ad-supported product at some price P_a , which may be zero (i.e., free). If the monopolist offers the ad-free software, its profit function is given by $\Pi_0 = (1 - P_0)P_0$, which is maximized at $P_0^* = 1/2$ yielding a maximum profit of $\Pi_0^* = 1/4$.




Figure 4.1: User segment for the ad-supported software

If instead the monopolist offers an ad-supported product, a consumer of type (μ, θ) will use the software if and only if $u_a(\mu, \theta, P_a) \geq 0$, i.e., if $\mu \geq P_a + \theta$. Figure 4.1 illustrates which consumers will use the ad-supported software; depending on P_a and λ , the user segment is represented either by a triangular or a trapezoidal area. This necessitates the consideration of different cases for λ when we determine the monopolist's optimal ad-supported pricing strategy. Table 4.2 lists the monopolist's optimal price and profit when it offers only the ad-supported software. As we intuitively expect, the optimal price for the ad-supported software is a non-increasing function of the per-user advertising rate (r), and the optimal profit is a monotonically increasing function of r and a monotonically decreasing function of λ . Using results in Table 4.2, we can easily prove Proposition 1. All proofs are given in the Appendix of this chapter.

Table 4.2: Optimal price and profit of the ad-supported software ³⁴

λ range	r range	Optimal price (P_a^*)	Optimal profit (Π_a^*)
$\lambda \in (0, \frac{2}{3}]$	$r \in (0, \frac{2-\lambda}{2}]$	$\frac{2-\lambda-2r}{4}$	$\frac{(2-\lambda+2r)^2}{16}$
	$r \in [\frac{2-\lambda}{2}, \infty)$	0	$\frac{(2-\lambda)r}{2}$
$\lambda \in [\frac{2}{3}, 1]$	$r \in (0, \frac{3\lambda-2}{3}]$	$\frac{1-2r}{3}$	$\frac{2(1+r)^3}{27\lambda}$
	$r \in [\frac{3\lambda-2}{3}, \frac{2-\lambda}{2}]$	$\frac{2-\lambda-2r}{4}$	$\frac{(2-\lambda+2r)^2}{16}$
	$r \in [\frac{2-\lambda}{2}, \infty)$	0	$\frac{(2-\lambda)r}{2}$
$\lambda \in [1, \infty)$	$r \in (0, \frac{1}{2}]$	$\frac{1-2r}{3}$	$\frac{2(1+r)^3}{27\lambda}$
	$r \in [\frac{1}{2}, \infty)$	0	$\frac{r}{2\lambda}$

³⁴ Proofs for all these cases are provided in the proof of Proposition 1 in the Appendix of this chapter.

Proposition 1: *Suppose that the monopolist can offer its application only on one platform—either ad-free or ad-supported. If $r > \frac{\lambda}{2}$, the monopolist will prefer offering ad-supported software (at a reduced price) to offering ad-free software; if $r < \frac{\lambda}{2}$, it will prefer offering the ad-free software rather than the ad-supported software. If $r > \max[\frac{2-\lambda}{2}, \frac{1}{2}]$, the monopolist will set a zero price for its ad-supported software, i.e., it will rely solely on advertising revenues.*

Though we do not model how firms might be able to influence λ and r , we do note that these values may differ significantly among firms across industries. For example, if a firm's software and advertising displays are visually pleasing and relevant to users, we expect the firm to have a relatively low λ or a relatively high r . Google is a leader in providing free online ad-supported applications earning the bulk of its revenues from advertising. In our framework, Google's strategy is a result of its high r value, which is a manifestation of its ability to do well-targeted, contextual advertising (e.g., based on keywords used in users' emails or search strings). Many software firms offer only ad-free software; this may be because of their high λ values (i.e., many consumers are very bothered by ads displayed via their software) for such types of applications or because of their low r values indicating their low ability to do effective advertising. While in practice many firms offer only ad-free or only ad-supported software, some firms do offer both versions of their applications; they may sell the ad-free version for a fee and also offer the ad-supported version either at a reduced price or for free. We next analyze the concurrent offering of both ad-free and ad-supported software.

4.2.2. Both ad-supported and ad-free software

If the firm offers both ad-free and ad-supported software, it will face some cannibalization between its two offerings.³⁵ Some consumers who would otherwise pay for the ad-free software will

³⁵ If the price of the software is so high that no consumers use it, then that product is effectively not offered by the firm. In this paper, we say that a product is offered only if it has a non-zero market share.

switch to the cheaper (or free) ad-supported software. Consumers of type (μ, θ) will use the ad-supported software if $u_a(\mu, \theta, P_a) \geq 0$ and $u_a(\mu, \theta, P_a) > u_0(\mu, P_0)$, and will buy the ad-free software if $u_0(\mu, P_0) \geq 0$ and $u_0(\mu, P_0) \geq u_a(\mu, \theta, P_a)$. That is, all consumers with $\mu \geq P_a + \theta$ and $\theta < P_0 - P_a$ will use the ad-supported software, and those with $\mu \geq P_0$ and $\theta \geq P_0 - P_a$ will buy the ad-free software. Note that if both types of software have non-zero market shares, we must have $P_0 > P_a$ (otherwise no one will use the ad-supported software) and $P_0 - P_a < \lambda$ (otherwise no one will buy the ad-free software). Figure 4.2 illustrates how consumers are segmented in the market.

We can easily express the firm's profit function as the sum of two parts—one from the ad-free version and one from the ad-supported version—as given by (1).

$$\Pi_{both} = \frac{1}{\lambda}(\lambda - P_0 + P_a)(1 - P_0)P_0 + \frac{1}{2\lambda}(2 - P_0 - P_a)(P_0 - P_a)(P_a + r) \quad (1)$$

Propositions 2 and 3 can be proved with the aid of *Lemma 1*.

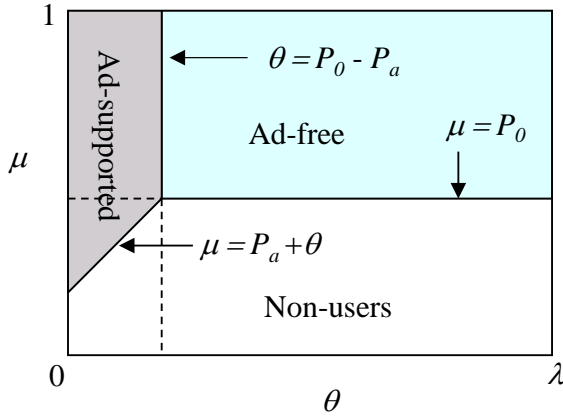


Figure 4.2: user segments for ad-free and ad-supported software

Lemma 1: For any $\lambda > 0$ and $r > 0$, the monopolist will make strictly higher profits offering both ad-free and ad-supported software than offering only ad-free software.

Proposition 2: For any $\lambda > 1$ and $r > 0$, the monopolist will make strictly higher profits offering both ad-free and ad-supported software than offering either one alone.

Proposition 3: For any $\lambda \leq 1$, there exists some r^* such that for all $r < r^*$, the monopolist will find it optimal to offer both ad-free and ad-supported software, and for all $r > r^*$, the monopolist will find it optimal to offer only the ad-supported software. In addition, if $r > \max[\frac{2-\lambda}{2}, r^*]$, the monopolist will offer only its ad-supported software and let consumers use it free of charge.

Propositions 2 and 3 explain why some firms find it optimal to offer both ad-free and ad-supported software while others may offer only an ad-supported version even when all fixed costs are neglected. Many applications such as peer-to-peer software (e.g., Kazaa) and instant messengers (e.g., MSN Messenger) are offered only as a free ad-supported product. Most consumers do not seem to be bothered much by the advertisements in such applications. Software firms or their third-party advertising agency can potentially utilize user profile information and usage contexts to display more targeted ads. This implies a relatively low λ value and a relatively high r , making it optimal for firms to offer only ad-supported software. Some consumer software applications have both an ad-supported version and an ad-free version. For example, Winamp, a multimedia player made by Nullsoft, has both a free ad-supported version and a paid, ad-free version.³⁶ Our analysis suggests that for such applications, many users may have high distaste for advertisements or the per-user advertising rate is not very high. Earlier researchers (Jones and Mendelson 1998, Bhargava and Choudhary 2001) have shown that a monopolist finds versioning of quality differentiated (ad-free) software to be sub-optimal. Other researchers (Jing 2003, Wu and Chen 2007) have shown that versioning can be optimal if there is network externality or piracy in the market. Our analysis shows that versioning on ad-free and ad-supported platforms can also be optimal.

In practice, many consumer applications are still offered only as ad-free software even though ad-supported versions are technically feasible. For example, WinZip, a popular zip and extraction application, is offered only as paid ad-free software.³⁷ According to Propositions 2 and 3, offering only ad-free software should be sub-optimal. But recall that our analysis assumes that the related fixed cost is sunk. If

³⁶ <http://www.winamp.com/>, accessed in June 2007.

³⁷ <http://www.winzip.com/>, accessed in June 2007.

the fixed cost of developing an ad-supported product is not sunk, firms must compare this fixed cost with the potential gain in advertising revenue. Many firms offering only their existing ad-free software likely perceive it to be too costly (in terms of the required fixed cost and product cannibalization) to develop an ad-supported version for the gain in advertising fees. However, these firms may later decide to introduce an ad-supported version if their expected advertising rate increases or if the required fixed cost decreases (e.g., due to their improved technology or capability). Next we extend our analysis to a competitive setting.

4.3. Competitive Market

4.3.1. Vertically differentiated market with only ad-free software (Baseline case)

We now consider two firms each of which offers its ad-free software competing in a vertically differentiated market. This competitive scenario will serve as a benchmark in later sub-sections. Firms' products are of different quality levels. A consumer of type μ derives a net utility of $u_i(\mu, P_i) = q_i\mu - P_i$ from product i (of firm i), where P_i and q_i are product i 's price and quality, respectively, and $\mu \sim \text{uniform}[0,1]$. Such linear utility functions are commonly used in the literature (e.g., Mussa and Rosen 1978, Shaked and Sutton 1982, Moorthy 1988). Without loss of generality, we assume that product 2's quality is at least as high as that of product 1, i.e., $q_2 \geq q_1$. Each consumer buys at most one product and derives zero utility if s/he does not purchase any product. A consumer of type μ will buy product i if and only if $u_i(\mu, P_i) \geq 0$ and $u_i(\mu, P_i) > u_j(\mu, P_j)$ for $j \neq i$. These two conditions are commonly known as the individual rationality and incentive compatibility constraint, respectively. Firms maximize their profits by strategically setting their prices. Lemma 2 shows the market equilibrium prices and profits.

Lemma 2: Suppose that firms have zero marginal costs and that their fixed costs are sunk. Then, the

equilibrium prices are given by $P_1^* = \frac{(q_2 - q_1) \cdot q_1}{4q_2 - q_1}$ and $P_2^* = \frac{2(q_2 - q_1) \cdot q_2}{4q_2 - q_1}$; firms' equilibrium

profits are $\pi_1^* = \frac{(q_2 - q_1) \cdot q_1 q_2}{(4q_2 - q_1)^2}$ and $\pi_2^* = \frac{4(q_2 - q_1) \cdot q_2^2}{(4q_2 - q_1)^2}$, respectively.

4.3.2. Vertically differentiated market with both ad-free and ad-supported software

We now analyze a vertically differentiated duopoly market in which one firm offers ad-free software while the other offers free ad-supported software. Firm 1's and firm 2's software quality levels are denoted by q_1 and q_2 , respectively. Without loss of generality, we normalize the quality of the high quality product to 1; i.e., $q_2 = 1$ and $q_1 < 1$. A consumer of type (μ, θ) derives a net utility of $u_i = q_i \mu - \theta$ from the free ad-supported software of quality q_i , and $u_j = q_j \mu - P_j$ from the ad-free software of quality q_j , which is priced at P_j . As before, consumers' willingness to pay for quality and distaste for ads are uniformly distributed: $\mu \sim \text{uniform}[0,1]$ and $\theta \sim \text{uniform}[0,\lambda]$, where $\lambda > 0$.

We analyze two scenarios—firm 1 offers ad-supported software while firm 2 offers ad-free software, and firm 1 offers ad-free while firm 2 offers ad-supported. We first consider the scenario under which firm 1 (the low quality firm) offers free ad-supported software while firm 2 (the high quality firm) sells ad-free software. Individual rationality and incentive compatibility constraints dictate that a consumer of type (μ, θ) will use the free ad-supported software if $u_1 \geq 0$ and $u_1 > u_2$, and will buy the ad-free software if $u_2 \geq 0$ and $u_2 \geq u_1$. It is easy to show that all consumers with $\mu \geq \frac{\theta}{q_1}$ and

$\mu < \frac{P_2 - \theta}{1 - q_1}$ will use the ad-supported software, those with $\mu \geq \max[P_2, \frac{P_2 - \theta}{1 - q_1}]$ will buy the ad-free

software, and the rest will use neither. Figure 4.3 illustrates how consumers are segmented in two possible

cases based on whether $P_2 > 1 - q_1$ or $P_2 \leq 1 - q_1$. The profit function of the high quality firm can be simplified to (2).




Figure 4.3: Vertically differentiated market: ad-free ($q_2=1$) and ad-supported ($q_1<1$)

$$\Pi_2 = \begin{cases} \frac{1}{2\lambda} P_2 (1 - P_2) [2\lambda + 1 - q_1 - (1 + q_1) P_2], & \text{if } P_2 \geq 1 - q_1; \\ \frac{1}{\lambda} P_2 \left[\lambda(1 - P_2) - \frac{q_1^2 P_2^2}{2(1 + q_1)} \right], & \text{if } P_2 \leq 1 - q_1 \end{cases} \quad (2)$$

Note that Π_2 is a continuous function of P_2 on the (practically meaningful) closed interval $[0, \min[\lambda, 1]]$ and that, in both price ranges specified in (2), it is a cubic polynomial in terms of P_2 . Further, since $\Pi_2 = 0$ on both end points and $\Pi_2 > 0$ inside the interval, there exists some P_2^* in the interval that will yield a positive, maximum profit. The low quality firm's profit can easily be expressed in terms of P_2^* as given by (3).

$$\Pi_1 = \begin{cases} \frac{r}{2\lambda} \left[P_2^{*2} q_1 + (P_2^* q_1 + P_2^* - 1 + q_1)(1 - P_2^*) \right], & \text{if } P_2^* \geq 1 - q_1; \\ \frac{P_2^{*2} q_1 r}{2\lambda(1 - q_1)}, & \text{if } P_2^* \leq 1 - q_1 \end{cases} \quad (3)$$




Figure 4.4: Vertically differentiated market: ad-supported ($q_2=1$) and ad-free ($q_1<1$)

We now consider the second scenario, under which firm 2 offers free ad-supported software while firm 1 sells ad-free software. Our analysis is similar to the preceding one. Figure 4.4 illustrates how consumers are segmented in two possible cases based on whether $P_1 \leq \lambda - 1 + q_1$ or $P_1 > \lambda - 1 + q_1$.

Consumers of type (μ, θ) with $\frac{P_1}{q_1} \leq \mu < \frac{\theta - P_1}{q_2 - q_1}$ will buy the low quality ad-free software while those

with $\mu \geq \max[\theta, \frac{\theta - P_1}{q_2 - q_1}]$ will use the free ad-supported software. The profit function of the high

quality firm can be simplified to (4).

$$\Pi_1 = \begin{cases} \frac{P_1}{2\lambda} \left(\lambda - \frac{P_1}{q_1} \right) \left(\frac{\lambda - P_1}{1 - q_1} - \frac{P_1}{q_1} \right), & \text{if } P_1 > \lambda - 1 + q_1; \\ \frac{P_1}{2\lambda} \left(1 - \frac{P_1}{q_1} \right) \left(2\lambda - 1 + q_1 - P_1 - \frac{P_1}{q_1} \right), & \text{if } P_1 \leq \lambda - 1 + q_1 \end{cases} \quad (4)$$

Valid prices fall on the closed interval $[0, \min[q_1, \lambda q_1]]$, outside which firm 1 will earn a negative or zero profit. Π_1 is a continuous function of P_1 on this interval, and $\Pi_1 = 0$ on both end points and $\Pi_1 > 0$ inside the interval. Thus, there exists some P_1^* in the interval that will yield a positive, maximum profit.

Firm 2's profit can then be expressed by (5).

$$\Pi_2 = \begin{cases} \frac{r}{2\lambda} \left[\left(2 - \frac{P_1^*}{q_1} \right) \frac{P_1^*}{q_1} + \left(\lambda - \frac{P_1^*}{q_1} \right) \left(2 - \frac{P_1^*}{q_1} - \frac{\lambda - P_1^*}{1 - q_1} \right) \right], & \text{if } P_1^* > \lambda - 1 + q_1; \\ \frac{r}{2\lambda} \left[\left(2 - \frac{P_1^*}{q_1} \right) \frac{P_1^*}{q_1} + \left(1 - \frac{P_1^*}{q_1} \right) \left(1 - q_1 + P_1^* - \frac{P_1^*}{q_1} \right) \right], & \text{if } P_1^* \leq \lambda - 1 + q_1 \end{cases} \quad (5)$$

Comparing the above two competitive scenarios with that in sub-section 4.3.1, we prove Proposition 4 to show that unless q_1 and q_2 are very far apart, both firms are better off if either firm chooses to offer ad-supported software than if neither does.

Proposition 4: *For any $\lambda > 0$ and $r > 0$, there exists some $q^* \in [0, 1]$ such that for all $q_1 \in (q^*, 1]$, both firms will be better off if either firm offers its software as a free ad-supported product than if neither does.*

To illustrate Proposition 4, we take a concrete example of $\lambda = 1$ and $r = 0.1$. Suppose that firm 1 offers free ad-supported software while firm 2 sells ad-free software. We plot the equilibrium prices and profits in terms of q_1 for both this competitive scenario and the one in sub-section 4.3.1. As indicated in Figures 4.5 to 4.7, the dark (black) curves correspond to the case in which both firms sell ad-free software, and the light (blue) curves correspond to the case in which firm 1 offers ad-supported software while firm 2 sells ad-free software. Figure 4.5 shows that in this example, regardless of q_1 , firm 2 is always better off if firm 1 switches from ad-free to ad-supported software. Intuitively, this is because when firm 1 adopts the ad-supported platform to rely solely on advertising revenues, the price competition between the two firms is alleviated compared with when both sell ad-free products. This reduced competitive pressure allows firm 2 to charge much higher equilibrium prices as shown in Figure 4.6. Further, Figure 4.7 shows that, unless q_1 is very low, firm 1 also makes a higher profit offering ad-supported software than offering ad-free software. The results and the intuition are similar for the comparison with the competitive case in which firm 2 offers ad-supported software while firm 1 offers ad-free software.

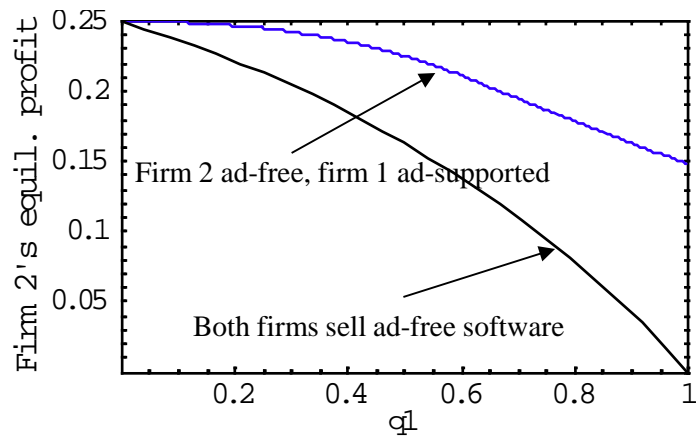


Figure 4.5: High quality firm's equilibrium profits in terms of q_1

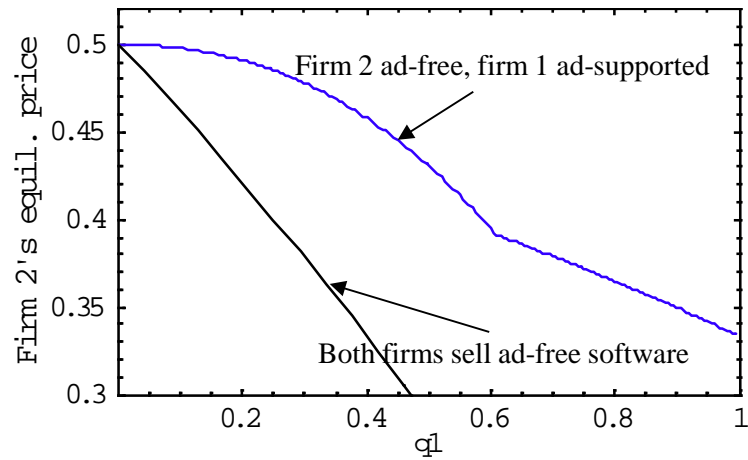


Figure 4.6: High quality firm's equilibrium prices in terms of q_1

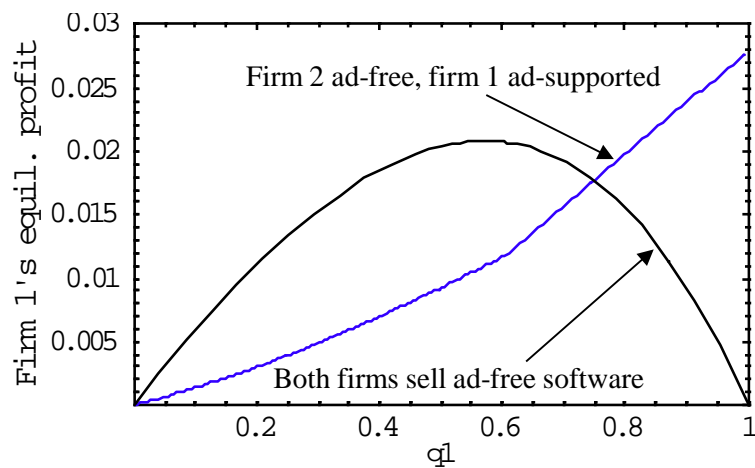


Figure 4.7: Low quality firm's equilibrium profits in terms of q_1

4.3.3. Competitive endogenous decisions to adopt ad-supported platform

We now extend our model to study firms' decisions of whether to sell their software ad-free or to offer it on an ad-supported platform. In the current game, each firm will first decide on the software platform—whether to offer its application ad-free or to distribute it as free, ad-supported software. In the second stage, the firm(s) that chose to offer ad-free software will set the price(s). In the last stage, consumers make decisions on software use—purchase an ad-free product, use free ad-supported software, or neither—and subsequently firms realize their profits. We adopt the same notations as used before. The analytical framework is summarized below. All parameters are common knowledge among firms and consumers. Firm 2 has a higher quality (q_2) and firm 1 has a lower quality ($q_1 < q_2$); we assume that the product quality levels are given (with fixed costs sunk).³⁸ A firm earns a per-user advertising rate of r from its ad-supported software. Consumers are heterogeneous in terms of willingness to pay for quality and distaste for ads: $\mu \sim \text{uniform}[0,1]$ and $\theta \sim \text{uniform}[0,\lambda]$, where $\lambda > 0$. A consumer of type (μ, θ) derives a net utility of $u = q\mu - \theta$ from the free ad-supported software (of quality q), and $u = q\mu - P$ from the ad-free software (of quality q at price P). We need to analyze four possibilities for firms' platform decisions since either firm may adopt either the ad-free or the ad-supported platform. Proposition 5 follows from the analyses of these cases.

Proposition 5: *Given λ , q_1 and q_2 , there exists some $r^* > 0$ such that if $r > r^*$, at equilibrium, the high quality firm will offer free ad-supported software and the low quality firm will sell ad-free software; further, if $r < r^*$, the high quality firm will offer ad-free software and there exists some $r' > 0$ such that the low quality firm will offer ad-free software if $r < r'$ and ad-supported software if $r > r'$.*

In essence, the high quality firm plays a dominant role in software platform decisions. If the advertising rate (r) is relatively high compared with an average consumer's willingness to pay for software, the high quality firm will offer ad-supported software, which forces the low quality firm to

³⁸ To allow for fully endogenous product quality decisions, we need to introduce another stage to the game, which is out of scope of this research. However, we will later allow firms to vary the quality of ad-supported software.

abandon that platform. In a vertically differentiated market, no consumers will use the low quality ad-supported software if the high quality software is also a free ad-supported product. Hence, in a vertically differentiated market, if online advertising is very lucrative, at equilibrium, the high quality firm will choose the ad-supported platform whereas the low-quality firm will offer an ad-free product to target the segment of consumers with high distaste for ads. But if r is not very high (or consumers' distaste for ads is generally high), the high quality firm will offer ad-free software. In that case, the low quality firm will also sell ad-free software if r is very low, or it will offer ad-supported software if r is in the middle range.

4.3.4. Competitive multi-platform decisions

We now investigate how the market outcome may change if each competitive firm may adopt *either* or *both* platforms. That is, both the high quality and the low quality firm can choose to offer ad-free software, free ad-supported software, or both. In addition, firms may potentially reduce the quality of their free ad-supported software (from their ad-free version). The current competitive game has three stages. In the first stage, firms simultaneously decide on which software platform(s) to adopt and the quality levels of their ad-supported software (if any). In the second stage, firms simultaneously set software prices. In the last stage, consumers make decisions on which software to use or purchase and firms thereby realize their profits.

We assume that firms' software development costs are sunk and that they incur no additional cost to reduce the quality of their ad-supported version (e.g., by removing some features). We adopt the same notations as used before. Firm 2 has a higher quality (q_2) and firm 1 has a lower quality ($q_1 < q_2$). Firm 1 may offer free ad-supported software at any quality $q_1' \leq q_1$, and firm 2 may offer free ad-supported software at any quality $q_2' \leq q_2$. Firms can obtain a per user advertising rate of r from their ad-supported software. As before, a consumer of type (μ, θ) derives a net utility of $u = q\mu - \theta$ from the free ad-supported software (of quality q), and $u = q\mu - P$ from the ad-free software (of quality q at price P). Consumers are uniformly distributed: $\mu \sim \text{uniform}[0,1]$ and $\theta \sim \text{uniform}[0,\lambda]$, where $\lambda > 1$. In the

current setting, both firms can adopt either or both software platforms. Analyses of all possible scenarios yield the following proposition.

Proposition 6: *In the competitive game described above, at equilibrium,*

- (a) *The low quality firm (firm 1) will offer only ad-free software;*³⁹
- (b) *If r is above some threshold, firm 2 will offer both ad-free and ad-supported software, the latter of which has a quality no smaller than q_1 .*

We find that, if the ad-supported platform is very profitable (as indicated by a large r value), firm 2 will offer both types of software. The quality of its ad-supported version depends on how large r is. For very large values of r , firm 2 will offer its ad-supported version at the same quality (q_2) as its ad-free software; but as r decreases, firm 2 will find it optimal to gradually reduce the quality of its ad-supported version from q_2 down to as low as q_1 . This is somewhat intuitive in that if r is large, firm 2 will want to attract a large number of consumers to use its ad-supported software and hence will provide high quality ad-supported software. If r is not large, a very high quality ad-supported offering will significantly cannibalize firm 2's own ad-free offering; firm 2 will thus have incentives to reduce cannibalization by lowering the quality of its ad-supported software. It seems that firm 2 has much to gain with the ad-supported platform if r is large. What if r is very small, say, as r approaches or equals zero? One will intuitively expect that, if $r = 0$, firm 2 can no longer derive *any* benefit from the new ad-supported platform. We will show that this is in fact *not* the case. In addition, naïve intuition tells us that firm 1 has much to lose in the presence of the new ad-supported platform; after all, firm 2 will have both ad-free and ad-supported software competing with firm 1's only ad-free product. We find that this intuition does *not* necessarily hold true, either.

³⁹ We say that "firm 1 offers only ad-free software" if it earns zero advertising revenue (as in the case in which firm 1 offers ad-supported software at any quality lower than that of firm 2's ad-supported software).




Figure 4.8: Competitive multi-platform decision

From Proposition 6, we know that firm 1 earns zero advertising revenue at equilibrium while firm 2 may offer ad-supported software at some quality $q \in [q_1, q_2]$ in addition to its ad-free version. Figure 4.8 illustrates one possible scenario of how consumers are segmented in this market. From individual rationality and incentive compatibility constraints, it is straightforward (albeit a little tedious) to derive firms' profits in terms of prices and qualities.

$$\Pi_1 = \frac{P_1}{2\lambda} \left(\frac{P_2 - P_1}{q_2 - q_1} - \frac{P_1}{q_1} \right) \left[2\lambda - \frac{(q + q_1)P_1}{q_1} - \frac{(q - q_1)(P_2 - P_1)}{q_2 - q_1} \right].$$

$$\text{If } P_2 \leq q_2 - q_1, \Pi_2 = \frac{P_2}{\lambda} \left[\lambda \left(1 - \frac{P_2 - P_1}{q_2 - q_1} \right) - \frac{1}{2} \left(\frac{P_2}{q_2 - q} - \frac{P_2 - P_1}{q_2 - q_1} \right) \left(P_1 + \frac{(q - q_1)(P_2 - P_1)}{q_2 - q_1} \right) \right] + \frac{r}{2\lambda} \times$$

$$\left[\left(\frac{P_2}{q_2 - q} - \frac{P_2 - P_1}{q_2 - q_1} \right) \left(P_1 + \frac{(q - q_1)(P_2 - P_1)}{q_2 - q_1} \right) + \left(\frac{(q - q_1)(P_2 - P_1)}{q_2 - q_1} + \frac{(q + q_1)P_1}{q_1} \right) \left(\frac{P_2 - P_1}{q_2 - q_1} - \frac{P_1}{q_1} \right) + \frac{P_1^2 q}{q_1^2} \right].$$

$$\text{If } P_2 \geq q_2 - q_1, \Pi_2 = \frac{P_2}{2\lambda} \left(2\lambda - P_2 + q_2 - q - P_1 - \frac{(q - q_1)(P_2 - P_1)}{q_2 - q_1} \right) \left(1 - \frac{P_2 - P_1}{q_2 - q_1} \right) + \frac{r}{2\lambda} \times$$

$$\left[\left(1 - \frac{P_2 - P_1}{q_2 - q_1} \right) \left(P_2 - q_2 + q + P_1 + \frac{(q - q_1)(P_2 - P_1)}{q_2 - q_1} \right) + \left(\frac{(q - q_1)(P_2 - P_1)}{q_2 - q_1} + \frac{(q + q_1)P_1}{q_1} \right) \left(\frac{P_2 - P_1}{q_2 - q_1} - \frac{P_1}{q_1} \right) + \frac{P_1^2 q}{q_1^2} \right]$$

Note that firm 2's profit function has two different forms depending on whether $P_2 \geq q_2 - q_1$. Firms will maximize their own profits simultaneously with firm 1 choosing P_1 and firm 2 choosing both P_2 and q (the quality of its ad-supported software). Though it is extremely difficult (even if possible) to obtain closed form solutions, it is rather easy to numerically compute the equilibrium profits and the corresponding equilibrium decision variable values for each given parameter value. We have systematically computed the equilibria for a large number of combinations of parameter values over wide ranges. We then compare the current equilibrium outcomes with those in the baseline case (in sub-section 4.3.1). Our main numerical results are summarized in the following propositions.

Proposition 7: *Unless both r and q_1 are very small, firm 2 will offer both ad-supported and ad-free software; even if $r = 0$, firm 2 will still benefit from offering ad-supported software if the two firms' quality levels (q_1 and q_2) are relatively close.⁴⁰*

Proposition 8: *If r is above some threshold (for any given q_1) or if q_1 is very close to q_2 (for any given r), firm 1 will benefit from the availability of the ad-supported platform even though it earns no advertising revenue at equilibrium. In particular, when the two firms' quality levels (q_1 and q_2) are relatively close, even if $r=0$, both firms will earn higher equilibrium profits with the availability of the ad-supported platform than without.*

Though it is not unexpected that the high quality firm will benefit from the ad-supported platform when the advertising rate (r) is large, it is surprising that the firm can benefit from offering ad-supported software *even if* $r = 0$. The underlying reason is the following. Note that in the extreme case of $r = 0$, firm 2 makes no profit from customers who use its ad-supported software (at quality q_1). However, the offering of the ad-supported software changes both firms' payoff functions in such a way that the equilibrium prices are both higher than if the ad-supported software is not offered. The reduced price competition

⁴⁰ Interestingly, in this extreme case of $r = 0$, the Nash equilibrium is that both firms offer their own ad-free software and both also offer ad-supported software at quality q_1 .

between the two firms explains why both firms may potentially benefit from the ad-supported platform even when $r = 0$. This effect of reduced price competition is somewhat counter-intuitive in that we generally expect the competition to be higher when more products are available. It is easier to see the “correct” intuition for the case of a large r value. If r is large, firm 2 will benefit greatly if more consumers use its ad-supported software; this gives firm 2 an incentive to increase the price of its ad-free software to make some customers switch to its ad-supported software, hence alleviating the price competition between firms’ ad-free software. Unless r is very close to zero *and* firm 1’s quality is very low, the net effect of ad-supported software on firm 2 will be an increase in its profit. The net effect of the availability of the ad-supported platform on firm 1’s profit can also be either positive or negative, because, on the one hand, firm 1 faces increased product competition (since firm 2 added ad-supported software), on the other hand, the price competition between the two ad-free products is lower. As Proposition 8 indicates, the net effect on firm 1’s profit (i.e., which of these two effects dominates) depends on r and the two firms’ relative quality levels, and it is possible for *both* firms to benefit from the availability of the ad-supported platform even if *neither* firm earns a positive advertising revenue or only one firm does. This suggests that in a quality differentiated software market where firms are intensely competing on prices, they may find it optimal to adopt the ad-supported platform even if they do not expect to make enough advertising revenue to cover the fixed cost required for the ad-supported software.

4.4. Conclusion

Enabled by Internet technologies, the newly emerged ad-supported licensing model has dramatically changed the software industries. The advantages of online advertising and the ubiquity of Internet access have made it more likely that an ad-supported software model will be viable in many application areas. Following the success of ad-supported software firms, many software makers are exploring the ad-supported platform in both consumer and business software markets. This research provides an economic analysis of the ad-supported software model, and fills the gap in the existing research literature on software licensing and also contributes to the growing two-sided market literature.

We have analyzed the economic impacts of the ad-supported platform in both monopoly and duopoly settings in a vertically differentiated market with heterogeneous consumers in terms of willingness to pay for quality and distaste for advertisements.

We show that, ignoring any fixed cost, it is generally sub-optimal for a monopolist to offer only ad-free software. Our analysis shows that versioning on ad-free and ad-supported platforms can in fact be optimal even though versioning on the ad-free platform alone is not. If some consumers' distaste for advertisements is high or if the per-user advertising rate is small, the monopolist will offer both ad-free and ad-supported software. However, if consumers' distaste for advertisements is relatively low compared with their valuation for the software features, the monopolist will offer only ad-supported software—either at a reduced price or for free—when its per-user advertising rate is high.

In a competitive, vertically differentiated market, in which each firm can adopt only one platform, we find that, first, if firms' quality levels are relatively close, both firms will be better off if either firm adopts the ad-supported platform than if neither does. Second, the high quality firm plays a dominant role in choosing the favorable platform. If the per-user advertising rate is above some threshold, the high quality firm will choose to offer free ad-supported software, forcing the low quality firm to abandon that platform to offer ad-free software. But if the per-user advertising rate is below that threshold, the high quality firm will offer ad-free software whereas the low quality firm may decide to adopt either platform depending on how large the advertising rate is.

When both competitive firms can potentially adopt multiple platforms, we find that, at equilibrium, the low quality firm offers only ad-free software, whereas the high quality firm will offer both ad-free and ad-supported software unless r is very close to zero *and* the low quality firm's product is very inferior. More interestingly, we find that even if neither firm earns a positive advertising revenue (i.e., $r = 0$) or only one firm does, both firms can still benefit from the availability of the ad-supported platform. We find that the ad-supported offering has a moderating effect on price competition between firms' ad-free offerings—an effect that may dominate the negative, cannibalization effect from the addition of more product offerings. This implies that, in a quality differentiated software market with

intense price competition, firms may have incentives to adopt the ad-supported platform even if they do not expect to make much advertising revenue to cover the incremental fixed cost.

As online advertising increases its significance and effectiveness in the digital economy, we expect more software firms to adopt the ad-supported platform either completely or as a complementary option. In future research, one may extend our framework to explicitly model the advertisers and examine firms' endogenous decisions of different types of advertising and the effectiveness of targeting software users. It may also be interesting to investigate how network effects influence firms' decisions and economic outcomes. This research has focused on a vertically differentiated market; future research can extend this analysis to a horizontally differentiated market. Analyzing a market with both horizontal and vertical product differentiation (e.g., in a way done in another context by Desai 2001) may also bring additional insights. Experimental studies on consumers' distastes for advertising in different types of software applications can also bring significant insight into the emergent ad-supported software phenomenon.

Appendix for Chapter 4

Proof of Proposition 1:

We first prove the results in Table 4.2. Note that we restrict the prices to be non-negative and that if the price is larger than 1, no consumers will use the software. Hence, the price interval of practical interest is $P_a \in [0, 1]$.

Case (1): $\lambda \in [1, \infty)$

The monopolist's profit is given by $\Pi_a(P_a) = \frac{(1-P_a)^2(P_a+r)}{2\lambda}$. Solving $\frac{d\Pi_a}{dP_a} = 0$, we obtain two critical points $P_a = \frac{1-2r}{3}$ and $P_a = 1$. $P_a = 1$ is clearly not optimal; $P_a = \frac{1-2r}{3}$ is in the valid price range only if $r \leq \frac{1}{2}$, in which case, we simply need to compare $\Pi_a(\frac{1-2r}{3})$ with $\Pi_a(0)$ to

find the optimal profit. It is straightforward to show that if $r \in (0, \frac{1}{2}]$, the monopolist's optimal price is

$$P_a^* = \frac{1-2r}{3} \text{ yielding a profit of } \Pi_a^* = \frac{2(1+r)^3}{27\lambda}, \text{ and that if } r \geq \frac{1}{2}, \text{ the optimal price is } P_a^* = 0$$

yielding a profit of $\Pi_a^* = \frac{r}{2\lambda}$.

Case (2): $\lambda \in (0, 1]$

Using Figure 4.1, we can write the profit function in two regions of P_a :

$$\Pi_a(P_a) = \begin{cases} \frac{(2-\lambda-2P_a)(P_a+r)}{2}, & \text{if } P_a \in [0, 1-\lambda]; \\ \frac{(1-P_a)^2(P_a+r)}{2\lambda}, & \text{if } P_a \in [1-\lambda, 1] \end{cases}$$

Solving $\frac{d\Pi_a}{dP_a} = 0$ using the above two functional forms of Π_a , we find three possible valid critical

points: $\frac{1-2r}{3}$, 1, and $\frac{2-\lambda-2r}{4}$. To be valid critical points, $P_a = \frac{1-2r}{3}$ should be on the interval

$[1-\lambda, 1]$ and $P_a = \frac{2-\lambda-2r}{4}$ should be on the interval $[0, 1-\lambda]$. $P_a = 1$ is clearly not optimal since

$\Pi_a(1) = 0$. To find the monopolist's optimal price and profit, we need to compare the profits at valid

critical points (if any) and at the boundaries ($P_a = 0$ and $P_a = 1-\lambda$). We discuss two intervals for λ : (i)

$\lambda \in (0, \frac{2}{3}]$ and (ii) $\lambda \in [\frac{2}{3}, 1]$.

Suppose $\lambda \in (0, \frac{2}{3}]$. Then, $P_a = \frac{1-2r}{3}$ is not a valid critical point because it is not inside the

interval $[1-\lambda, 1]$. Furthermore, one can easily to show that if $r \in (0, \frac{2-\lambda}{2}]$, the optimal price is

$P_a^* = \frac{2-\lambda-2r}{4}$ yielding a profit of $\Pi_a^* = \frac{(2-\lambda+2r)^2}{16}$, and if $r \in [\frac{2-\lambda}{2}, \infty)$, the optimal price is

$P_a^* = 0$ yielding a profit of $\Pi_a^* = \frac{(2-\lambda)r}{2}$.

Suppose $\lambda \in [\frac{2}{3}, 1]$. Then, if $r > \frac{2-\lambda}{2}$, neither $P_a = \frac{1-2r}{3}$ or $P_a = \frac{2-\lambda-2r}{4}$ is a valid

critical point, and the optimal price is thus $P_a^* = 0$ yielding a profit of $\Pi_a^* = \frac{(2-\lambda)r}{2}$. If $r \leq \frac{3\lambda-2}{3}$,

$P_a = \frac{1-2r}{3}$ is a valid critical point and $P_a = \frac{2-\lambda-2r}{4}$ is not; one then easily shows that the optimal

price is $P_a^* = \frac{1-2r}{3}$ yielding a profit of $\Pi_a^* = \frac{2(1+r)^3}{27\lambda}$. If $r \in [\frac{3\lambda-2}{3}, \frac{2-\lambda}{2}]$, $P_a = \frac{2-\lambda-2r}{4}$ is

a valid critical point and the monopolist's optimal price is $P_a^* = \frac{2-\lambda-2r}{4}$ yielding a profit of

$\Pi_a^* = \frac{(2-\lambda+2r)^2}{16}$. This completes the proof for all results in Table 4.2.

Recall that the monopolist's optimal profit for offering the ad-free software is $\Pi_0^* = 1/4$ at

$P_0^* = 1/2$. From Table 4.2, it is trivial to prove that $P_a^* < P_0^*$ and that if $r > \max[\frac{2-\lambda}{2}, \frac{1}{2}]$, the

monopolist will rely solely on advertising revenues (i.e., $P_a^* = 0$). Solving $\Pi_a^* = \Pi_0^* = \frac{1}{4}$ using the

results in Table 4.2, we get $r = \frac{\lambda}{2}$ for all cases. Note that Π_a^* is a monotonically increasing function of r

and a monotonically decreasing function of λ . Thus, Proposition 1 follows. \square

Proof of Lemma 1:

Suppose that the firm sets prices at $P_0 = P_a = \frac{1}{2}$. Then, no consumers will use the ad-supported

software; (1) simplifies to $\Pi_{both} = (1 - P_0)P_0$, which is exactly the profit function of offering only the

ad-free software. However, one can easily compute $\left. \frac{\partial \Pi_{both}}{\partial P_a} \right|_{P_0=\frac{1}{2}, P_a=\frac{1}{2}} = -\frac{r}{2\lambda}$, which implies that the

firm can improve its profit by reducing P_a . Hence, *Lemma 1* follows. \square

Proof of Proposition 2:

By *Lemma 1*, offering both types of software dominates offering only the ad-free version. Now, suppose that the firm sets its prices at $P_a = P_a^*$ as given in Table 4.2 and $P_0 = 1$. Clearly, $\Pi_{both}|_{P_0=1, P_a=P_a^*} = \Pi_a^*$ since no one buys the ad-free software. It is easy to show that

$$\left. \frac{\partial \Pi_{both}}{\partial P_0} \right|_{P_0=1, P_a=P_a^*} = -\frac{\lambda + P_a^* - 1}{\lambda} < 0 \text{ since } P_a^* \geq 0 \text{ and } \lambda > 1. \text{ This implies that there exists some } \delta > 0$$

such that $\Pi_{both}|_{P_0=1-\delta, P_a=P_a^*} > \Pi_a^*$. Thus, offering both types of software also strictly dominates offering only the ad-supported version. Thus, Proposition 2 follows. \square

Proof of Lemma 2:

Consumer purchasing decisions observe the individual rationality and incentive compatibility constraints. That is, a consumer of type μ will purchase product i if $u_i(\mu, P_i) \equiv q_i \mu - P_i \geq 0$ and $u_i(\mu, P_i) > u_j(\mu, P_j)$ for $j \neq i$. From these constraints, it is straightforward to show that the consumer

will buy product 1 if $\frac{P_1}{q_1} \leq \mu < \frac{P_2 - P_1}{q_2 - q_1}$, and product 2 if $\frac{P_2 - P_1}{q_2 - q_1} \leq \mu \leq 1$. Since consumers are

assumed to be uniformly distributed on $[0, 1]$, firms' profit functions are given by

$$\pi_1 = P_1 \cdot \left(\frac{P_2 - P_1}{q_2 - q_1} - \frac{P_1}{q_1} \right) \text{ and } \pi_2 = P_2 \cdot \left(1 - \frac{P_2 - P_1}{q_2 - q_1} \right), \text{ respectively.}$$

Each firm will select its own optimal price given the other's price. The first order conditions are

$$\frac{\partial \pi_1}{\partial P_1} = \left(\frac{P_2 - P_1}{q_2 - q_1} - \frac{P_1}{q_1} \right) - \frac{P_1}{q_2 - q_1} - \frac{P_1}{q_1} = 0 \quad \text{and} \quad \frac{\partial \pi_2}{\partial P_2} = 1 - \frac{P_2 - P_1}{q_2 - q_1} + \frac{-P_2}{q_2 - q_1} = 0. \quad \text{Solving these}$$

conditions simultaneously, we obtain the equilibrium prices (as the second order conditions for profit

maximization clearly hold): $P_1^* = \frac{(q_2 - q_1) \cdot q_1}{4q_2 - q_1}$ and $P_2^* = \frac{2(q_2 - q_1) \cdot q_2}{4q_2 - q_1}$. Equilibrium profits are

easily computed: $\pi_1^* = \frac{(q_2 - q_1) \cdot q_1 q_2}{(4q_2 - q_1)^2}$ and $\pi_2^* = \frac{4(q_2 - q_1) \cdot q_2^2}{(4q_2 - q_1)^2}$. □

Proof of Proposition 3:

Recall that firm's profit function is given by (1). Let Π_0^* denote the monopolist's optimal profit if it offers only ad-free software, Π_a^* denote the optimal profit if it offers only ad-supported software, and Π_{both}^* denote the optimal profit if it offers both types of software.

First, we show that there exists some $\delta > 0$ such that for all $r < \delta$, the monopolist will find it optimal to offer both ad-free and ad-supported software. By Lemma 1, for any $r > 0$, $\Pi_{both}^* > \Pi_0^*$. Note that as $r \rightarrow 0$, offering only the ad-supported software is sub-optimal to offering only ad-free software. Put mathematically, there exists some $\delta > 0$ such that for all $r < \delta$, $\Pi_0^* > \Pi_a^*$. Thus, $\Pi_{both}^* > \Pi_0^* > \Pi_a^*$ for all $r < \delta$. That is, for all $r < \delta$, the monopolist will find it optimal to offer both ad-free and ad-supported software.

Second, we prove that there exists some $\varepsilon > 0$ such that for all $r > \varepsilon$, it is optimal for the monopolist to offer only ad-supported software. We will prove this by contradiction. Suppose that it is *not* optimal for the monopolist to offer only ad-supported software. Then, offering both types of software (with both having a non-zero market share) must lead to a higher profit than offering only ad-supported software.

For both types of software to have non-zero market shares, the prices must satisfy $P_a^* < P_0^*$ and

$0 < P_0^* < \min[1, \lambda + P_a^*]$. We will discuss two cases: (i) $0 < P_0^* < \lambda + P_a^* \leq 1$, and (ii)

$0 < P_0^* < 1 < \lambda + P_a^*$.

Case (i): $0 < P_0^* < \lambda + P_a^* \leq 1$

$$\begin{aligned}
\left. \frac{\partial \Pi_{both}}{\partial P_0} \right|_{P_0^*, P_a^*} &= -\frac{P_0^*(1-P_0^*)}{\lambda} + \frac{(1-2P_0^*)(\lambda-P_0^*+P_a^*)}{\lambda} + \frac{(1-P_0^*)(P_a^*+r)}{\lambda} \\
&= \frac{(1-P_0^*)(P_a^*+r-P_0^*)}{\lambda} + \frac{(1-2P_0^*)(\lambda-P_0^*+P_a^*)}{\lambda} \\
&= \frac{(1-P_0^*)(P_a^*+r-P_0^*+\lambda-P_0^*+P_a^*)}{\lambda} - \frac{P_0^*(\lambda-P_0^*+P_a^*)}{\lambda} \\
&= \frac{(1-P_0^*)(2P_a^*+r-2P_0^*+\lambda)}{\lambda} - \frac{P_0^*(\lambda-P_0^*+P_a^*)}{\lambda}
\end{aligned}$$

Since $0 < P_0^* < \lambda + P_a^* \leq 1$, we know $0 < P_0^* < 1$ and $0 < \lambda + P_a^* - P_0^* \leq 1 - P_0^*$. Thus,

$$\begin{aligned}
\left. \frac{\partial \Pi_{both}}{\partial P_0} \right|_{P_0^*, P_a^*} &> \frac{(1-P_0^*)(2P_a^*+r-2P_0^*+\lambda)}{\lambda} - \frac{1 \cdot (1-P_0^*)}{\lambda} = \frac{(1-P_0^*)(2P_a^*+r-2P_0^*+\lambda-1)}{\lambda} \\
&= \frac{(1-P_0^*)[2(\lambda+P_a^*-P_0^*)+(r-\lambda-1)]}{\lambda} > 0, \text{ for all } r \geq \lambda+1.
\end{aligned}$$

This implies that the monopolist can increase its profit by slightly increasing P_0^* , which contradicts the fact that P_0^* is optimal.

Case (ii): $0 < P_0^* < 1 < \lambda + P_a^*$

We will examine how much the monopolist's profit will vary if we deviate from the supposed optimal price P_a^* while keeping P_0^* the same. In particular, we set $P_a = 1 - \lambda$ and compute

$$\Pi_{both} \Big|_{P_0=P_0^*, P_a=1-\lambda} - \Pi_{both} \Big|_{P_0=P_0^*, P_a=P_a^*}.$$

Straightforward (but a little tedious) algebraic manipulation yields:

$$\begin{aligned}
& \Pi_{both} \Big|_{P_0=P_0^*, P_a=1-\lambda} - \Pi_{both} \Big|_{P_0=P_0^*, P_a=P_a^*} \\
&= \frac{1}{\lambda} \left\{ (1-P_0^*)(P_a^* + \lambda - 1)(1 - \lambda + r - P_0^*) \right\} \\
&+ \frac{1}{2\lambda} \left\{ \left[2(P_a^* + \lambda - 1)(P_0^* - P_a^*) + (P_a^* + \lambda - 1)^2 \right] (1 - \lambda + r) - (P_a^* + \lambda - 1)(P_0^* - P_a^*)(2 - P_0^* - P_a^*) \right\}
\end{aligned}$$

For all $r \geq \lambda + 1$, both terms are clearly positive, noting that $\lambda + P_a^* - 1 > 0$, $1 - P_0^* > 0$, and $P_0^* - P_a^* > 0$. This implies that the monopolist can increase its profit by deviating from P_a^* , which contradicts the fact that P_a^* is optimal.

Thus, in both cases, we have shown that it is optimal for the monopolist to offer only ad-supported software if $r \geq \lambda + 1$. Thereby, we have proved, by construction, that there exists some $\varepsilon > 0$ such that for all $r > \varepsilon$, it is optimal for the monopolist to offer only ad-supported software.

Clearly, if it is optimal for the monopolist to offer only ad-supported software at r , then it must also be optimal to offer only ad-supported software at any $r' > r$.

The thresholds must be the same for two claims, since the optimal profit function must be an increasing function in r . Note that the optimal profit from ad-supported software is a strictly increasing function in r . Thus, from the two claims that we have proved, we conclude that there exists some r^* such that for all $r < r^*$, the monopolist will find it optimal to offer both ad-free and ad-supported software, and for all $r > r^*$, the monopolist will find it optimal to offer only the ad-supported software.

In addition, from Table 4.2, we know that when only the ad-supported software is offered, the optimal price is zero if $r > \frac{2-\lambda}{2}$ (for the present case of $\lambda \leq 1$). Thus, if $r > \max[\frac{2-\lambda}{2}, r^*]$, the monopolist will find it optimal to offer only its ad-supported software and to give it to consumers free of charge. □

Proof of Proposition 4:

We set $q_2=1$ in the baseline case (Section 4.3.1). By *Lemma 2*, both firms' equilibrium profits (π_1^* and π_2^*) are continuous functions of q_1 on the closed interval $[0, 1]$ and both are zero at $q_1=1$. In the case in which the low-quality firm offers free ad-supported software while the high quality firm offers ad-free software, firms' equilibrium profits (Π_1^* and Π_2^*) are also continuous functions of q_1 on $[0, 1]$, and both are *strictly* positive on $(0,1]$. Note also that the equilibrium profit functions have a finite number of critical points (at which the derivatives with respect to q_1 are zero). Thus, there exists some $q^* \in [0,1)$ such that for all $q_1 \in (q^*, 1]$, both $\Pi_1^* > \pi_1^*$ and $\Pi_2^* > \pi_2^*$ hold, i.e., both firms will be better off if the low quality firm offers free ad-supported software while the high quality firm sells ad-free software than if both firms sell their ad-free software. \square

Proof of Proposition 5:

There are four possible equilibrium cases to consider.

- (i) Firm 2 offers ad-supported software; firm 1 also offers ad-supported software;
- (ii) Firm 2 offers ad-supported software; firm 1 offers ad-free software;
- (iii) Firm 2 offers ad-free software; firm 1 also offers ad-free software;
- (iv) Firm 2 offers ad-free software; firm 1 offers ad-supported software.

We will denote firms' equilibrium profits in each of these cases by $\Pi_1^{*(k)}$ and $\Pi_2^{*(k)}$, where the superscript k indicates which case it is.

In case (i), firm 1 makes zero profit $\Pi_1^{*(i)} = 0$ because no consumers would use the low quality ad-supported software when the high quality ad-supported software is free. In case (ii), firm 1 makes a positive profit $\Pi_1^{*(ii)} > 0$ because if it sets a price low enough, some consumers (with $\frac{P_1}{q_1} \leq \mu < \frac{\theta - P_1}{q_2 - q_1}$) will prefer buying the low quality ad-free software to using the free ad-supported software. Thus, at

equilibrium, if firm 2 offers ad-supported software, firm 1 will definitely choose to offer ad-free software rather than ad-supported software. Cases (iii) and (iv) have been analyzed in sections 4.3.1 and 4.3.2.

Now we consider firm 2's decision. For firm 2 to prefer offering ad-supported software, it must make a higher profit with ad-supported software, i.e., $\Pi_2^{*(ii)} > \Pi_2^{*(iii)}$ and $\Pi_2^{*(ii)} > \Pi_2^{*(iv)}$. Note that $\Pi_2^{*(ii)}$ is a linear function in r and that given λ , q_1 and q_2 , if firm 2 prefers offering ad-supported software at r , then it will also prefer ad-supported software at any $r' > r$. Hence, given λ , q_1 and q_2 , there exists some $r^* > 0$ such that if $r > r^*$, at equilibrium, firm 2 (the high quality firm) will offer free ad-supported software and firm 1 will sell ad-free software; further, if $r < r^*$, firm 2 will offer ad-free software.

We now examine firm 1's optimal decision when firm 2 has optimally decided to offer ad-free software (i.e., $r < r^*$). If firm 1 offers ad-supported software, its profit ($\Pi_1^{*(iv)}$) will be a linear function in r and approaches zero as r approaches zero. Note that $\Pi_1^{*(iii)} > 0$ (and is independent of r) and that given λ , q_1 and q_2 , if firm 1 prefers offering ad-free software to offering ad-supported software at r (i.e., $\Pi_1^{*(iii)} > \Pi_1^{*(iv)}$ at r), then it also prefers offering ad-free software at any smaller r . Hence, there exists some $r' > 0$ such that firm 1 (the low quality firm) will offer ad-free software if $r < r'$ and ad-supported software if $r > r'$. □

Proof of Proposition 6:

Suppose that, at equilibrium, firm 1 offers ad-supported software at some quality $q \leq q_1$ and makes some positive advertising revenue. This implies that firm 2 is not offering any ad-supported software at any quality higher than q , otherwise no consumers will use firm 1's lower quality ad-supported software. However, this outcome cannot be at equilibrium, because firm 2 will have an incentive to deviate from its current strategy; in particular, firm 2 can improve its profit if it also offers ad-supported software at a quality level infinitesimally higher than q . Thus, at equilibrium, firm 1 does not offer any ad-supported software. □




Figure 4A-1: effect of firm 2's adding ad-free product

We now examine firm 2's equilibrium strategy. From Proposition 5, it is easy to see that if r is large enough, firm 2 will offer an ad-supported product. Suppose that, at equilibrium, firm 2 offers only ad-supported software. We will show that firm 2 can always improve its profit by adding ad-free software (of quality q_2) to its ad-supported offering, and hence conclude that firm 2 will also offer ad-free software at equilibrium. As illustrated by the dashed lines in Figure 4A-1, if firm 2 adds its ad-free software to the market, it will lose some advertising revenue (the grey triangular area) and gain some software sales revenue (the upper rectangular area). The ratio of firm 2's sales gain to its advertising loss is computed below.

$$\frac{\text{Gain}}{\text{Loss}} = \frac{P_2 \left(1 - \frac{P_2 - P_1}{q_2 - q_1} \right) (\lambda - P_2)}{\frac{r}{2} \left(1 - \frac{P_2 - P_1}{q_2 - q_1} \right) \cdot \left(\frac{P_2 - P_1}{q_2 - q_1} - P_2 \right)} = \frac{2P_2(\lambda - P_2)}{r \left(\frac{P_2 - P_1}{q_2 - q_1} - P_2 \right)}$$

It is straightforward to show that the above ratio approaches infinity as $P_2 \rightarrow \frac{P_1}{q_1}$. That is, firm 2

can improve its profit by adding its ad-free software to the market at any price that is very close to this

limit price. Thus, at equilibrium, firm 2 must be offering its ad-free software. That is, we have shown that if r is large enough, firm 2 will offer both ad-free and ad-supported software at equilibrium. It is also clear that at, equilibrium, firm 2's ad-supported software must have a quality level no smaller than q_1 because otherwise firm 1 will have a profitable deviation from equilibrium by adding its own ad-supported software at a quality just higher than that of firm 2's ad-supported software. Finally, note that as $r \rightarrow 0$, firm 2 will offer only ad-free software. Hence, Proposition 6 follows. \square

References

- Adams, W. J., J. Yellen. 1977. What makes advertising profitable? *The Economic Journal* **87**(347) 427-449.
- Anderson, C. 2006. The long tail: Why the future of business is selling less of more. Hyperion: New York.
- Anderson, S., S. Coate. 2005. Market provision of broadcasting: A welfare analysis. *Review of Economic Studies* **72**(4) 947-972.
- Ansari, A., C.F. Mela, S.A. Neslin. 2008. Customer channel migration. *J. of Marketing Research* **45**(1) 60-76.
- Ante, S. "Amazon: Turning Consumer Opinions into Gold." *Business Week* 15 October 2009. Online: http://www.businessweek.com/magazine/content/09_43/b4152047039565.htm?chan=magazine+channel_special+report
- Araujo, A., D. Gottlieb, H. Moreira. 2008. A model of mixed signals with applications to countersignaling and the GEC. *RAND Journal of Economics* **38**(4) 1020-1043.
- Armstrong, M. 2006. Competition in two-sided markets. *RAND Journal of Economics* **37** (3) 668-691.
- Bagwell, K. 2007. The economic analysis of advertising. M. Armstrong, R. Porter, eds. *Handbook of Industrial Organization*, Vol. 3, Chapter 28. North-Holland, Amsterdam.
- Bakos, Y. 1997. "Reducing buyer search costs: Implications for electronic marketplaces." *Management Science* **43**(12) 1676-1692.
- Banks, J., J. Sobel. 1987. Equilibrium selection in signaling games. *Econometrica* **55**(3) 747-661.
- Baye, M., R. Gatti, P. Kattuman, J. Morgan. 2007. A dashboard for online pricing. *California Management Review* **50**(1) 202-216.
- Bhargava, H., V. Choudhary. 2001. Information goods and vertical differentiation. *Journal of Management Information Systems* **18**(2) 89-106.
- Biyalogorsky E., P. Naik. 2003. Clicks and mortar: The effect of online activities on offline sales. *Marketing Letters* **14**(1) 21-32.
- Brynjolfsson, E., Y. Hu, M. D. Smith. 2003. Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. *Management Science* **49**(11) 1580-1596.
- Brynjolfsson, E., Y. Hu, M. D. Smith. 2006. From niches to riches: The anatomy of the long tail. *Sloan Management Review* **47**(4) 67-71.

- Brynjolfsson, E., M. Smith. 2000. Frictionless commerce? A comparison of internet and conventional retailers. *Management Science* **46**(4) 563-585.
- Casadesus-Masanell, Ramon, Feng Zhu. 2009. Strategies to Fight Ad-sponsored Rivals. NET Institute Working Paper #09-09.
- Chen, P., S. Dhanasobhon, M. D. Smith. 2007. All reviews are not created equal: The disaggregate impact of reviews and reviewers at Amazon.com. *Proceedings of the International Conference on Information Systems*.
- Chen, Y., J. Xie. 2005. Third-party product review and firm marketing strategy. *Marketing Science* **23**(2) 218-240.
- Chen, Y. , J. Xie. 2008. Online consumer review: A new element of marketing communications mix. *Management Science* **54**(3) 477-491.
- Chevalier, J. A., D. Mayzlin. 2006. The effect of word of mouth on sales: Online book reviews. *Journal of Marketing Research* **43**(3), 345-354.
- Cho, I.K., D. Kreps. 1987. Signaling games and stable equilibria. *Quart. J. of Econ.* **102**(2) 179-222.
- Choi J., D. Bell. 2009. Preference minorities and the internet. . *J. of Marketing Research*. Forthcoming.
- Choi J., S. K. Hui, D. Bell. 2010. Spatio-temporal analysis of imitation behavior across new buyers at an online grocery retailer. *J. of Marketing Research* **47**(1) 65-79.
- Choudhary V., K. Tomak, A. Chaturvedi. 1998. Economic benefits of renting software. *Journal of Organizational Computing and Electronic Commerce* **8**(4) 277-305.
- Clemons, E. K., G. Gao, L. M. Hitt. 2006. When online reviews meet hyperdifferentiation: A study of the craft beer industry. *Journal of Management Information Systems* **23**(2) 149-171.
- comScore, The Kelsay Group. "Online Consumer-Generated Reviews Have Significant Impact on Offline Purchase Behavior." *comScore Press Release* 29 November 2007. Online: http://www.comscore.com/Press_Events/Press_Releases/2007/11/Online_Consumer_Reviews_Impact_Offline_Purchasing_Behavior
- Coughlan, A. T., B. Wernerfelt. 1989. On credible delegation by oligopolists: A discussion of distribution channel management. *Management Science* **35**(2) 226-239.
- D'Aspremont, C., J. J. Gabszewicz, J.-F. Thisse. 1979. On Hotelling's "Stability in competition." *Econometrica* **47**(5) 1145-1150.

- Dellarocas, C. 2006. Strategic manipulation of internet opinion forums: Implications for consumers and firms. *Management Science* **56**(10) 1577-1593.
- Dellarocas, C., N. F. Awad, X. Zhang. 2007. Exploring the value of online product ratings in forecasting sales: The case of motion pictures. *Journal of Interactive Marketing* **21**(4) 23-45.
- Desai, P. 2000. Multiple messages to retain retailers: Signaling new product demand. *Marketing Science* **19**(4) 381-389.
- Desai, P. S. 2001. Quality segmentation in spatial markets: When does cannibalization affect product line design? *Marketing Science* **20**(3) 265-283.
- Desai, P., O. Koenigsberg, D. Purohit. 2004. Strategic Decentralization and Channel Coordination. *Quantitative Marketing and Economics* **2**(1) 5-22.
- Desai, P. S., K. Srinivasan. 1995. Demand signaling under unobservable effort in franchising: Linear and non-linear price contracts. *Management Science* **41**(10) 1608-1623.
- Desiraju R., S. Moorthy. 1997. Managing a distribution channel under asymmetric information with performance requirements. *Management Science* **43**(12) 1628-1644.
- Dewan, R.M., M. Freimer, J. Zhang. 2003. Management and valuation of advertisement-supported Web sites. *Journal of Management Information Systems* **19**(3) 87-98.
- Duan, W., B. Gu, A. B. Whinston. 2008. The Dynamics of online WOM and product sales—An empirical investigation of the movie industry. *Journal of Retailing* **84**(2) 233-242.
- Economides, Nicholas, Evangelos Katsamakas. 2006. Two-sided competition of proprietary vs. open source technology platforms and the implications for the software industry. *Management Science* **52**(7) 1057-1071.
- Elberse, A. 2008. Should you invest in the long tail? *Harvard Business Review* **86**(7/8) 88-96.
- Eisenmann, T., G. Parker, M. W. Van Alstyne. 2006. Strategies for two-sided markets. *Harvard Business Review* October 2006.
- eMarketer. "What are UK Women Doing Online?" *The eMarketer Daily Newsletter* 7 April 2009. Online: <http://www.emarketer.com/Article.aspx?R=1007018>
- Fan, M., S. Kumar, A. Whinston. 2007. Selling or advertising: strategies for providing digital media online. *Journal of Management Information Systems* **24**(3) 143-166.
- Feltovich, N., R. Harbaugh, T. To. 2002. Too cool for school? Signalling and countersignalling. *RAND Journal of Economics* **33**(4) 630-649.

- Forman, C., A. Ghose, B. Wiesenfeld. 2008. Examining the relationship between reviews and sales: the role of social information in electronic markets. *Information Systems Research* **19**(3) 291-313.
- Gal-Or, Esther, Anthony J. Dukes. 2003. Minimum differentiation in commercial media markets. *Journal of Economics & Management Strategy* **12**(3) 291-325.
- Godes, D., D. Mayzlin. 2004. Using online conversations to measure word of mouth communication. *Marketing Science* **23**(4) 545-560.
- Holmstrom, B. 1979. Moral Hazard and Observability. *Bell Journal of Economics* **10**(1) 74-91.
- Huang, K., A. Sundararajan. 2005. Pricing models for on-demand computing. CeDER Working Paper No. 05-26. Available at SSRN: <http://ssrn.com/abstract=859504>.
- Iyer, G., M. Villas-Boas 2003. A bargaining theory of distribution channels. *J. of Marketing Research* **40**(1) 80-100.
- Jerath, K., Z. J. Zhang. 2010. Store Within a Store. *J. of Marketing Research* **47**(4) 748-763.
- Jeuland, A. P., S. M. Shugan. 1983. Managing channel profits. *Marketing Science* **2**(3) 239-272.
- Jiang, B., K. Srinivasan. 2011. Impact of online consumer reviews on competitive marketing strategies. *Carnegie Mellon University*. Working paper.
- Jiang, B., P. Chen, T. Mukhopadhyay. 2007. Software licensing: Pay-per-use versus perpetual. Working Paper. Available at SSRN: <http://ssrn.com/abstract=1088570>.
- Jiang, B. J., B. Wang. 2008. Impact of consumer reviews and ratings on sales, prices, and profits: theory and evidence. *Proceedings of the International Conference on Information Systems*.
- Jing, B. 2003. Market Segmentation for Information Goods with Network Externalities. Working Paper, Stern School of Business, New York University.
- Jones, R., H. Mendelson. 1998. Product and price competition for information goods. Working Paper, Graduate School of Business, Stanford University.
- Kornish, L. J. 2009. Are user reviews systematically manipulated? Evidence from the helpfulness ratings. Working Paper. University of Colorado at Boulder.
- Kuksov, D., Y. Xie. 2010. Pricing, frills, and customer ratings. *Marketing Science* **29**(5) 925-943.
- Lee, D., H. Mendelson. 2008. Divide and conquer: Competing with free technology under network effects. *Production and Operations Management* **17**(1) 12-28.
- Lal, R. and R. Staelin. 1986. Salesforce compensation plans in environments with asymmetric

- information. *Marketing Science* **5**(3) 179-198.
- Lewis, T., D. Sappington. 1994. Supplying information to facilitate price discrimination. *International Economic Review* **35**(2) 309–327.
- Li, X., L. M. Hitt. 2008. Self selection and information role of online product reviews. *Information Systems Research* **19**(4) 456-474.
- Liu, Y. 2006. Word of mouth for movies: its dynamics and impact on box office revenue. *Journal of Marketing* **70**(3) 74-89.
- Martin, J.M., G.M. Barron, M. I. Norton. 2008. Variance-seeking for positive (and variance-aversion for negative) experiences: risk-seeking in the domain of gains? Working Paper, Harvard Business School.
- Mayzlin, D. 2006. Promotional chat on the Internet. *Marketing Science* **25**(2) 155-163.
- Mayzlin, D., J. Shin. 2009. Uninformative advertising as an invitation to search. *Yale University*, Working paper.
- McGuire, T. W., R. Staelin. 1983. An industry equilibrium analysis of downstream vertical integration. *Marketing Science* **2**(2) 161-191.
- Moorthy, K. S. 1988. Product and price competition in a duopoly. *Marketing Sci.* **7**(2) 141-168.
- Moorthy, S., K. Srinivasan. 1995. Signaling quality with a money-back guarantee: The role of transaction costs. *Marketing Science* **14**(4) 442-466.
- Mussa, M., S. Rosen. 1978. Monopoly and product quality. *J. of Econom. Theory* **18**(2) 301-317.
- Neslin S. A., D. Grewal, R. Leghorn, V. Shankar, M. L. Teerling, J. S. Thomas, P. C. Verhoef. 2006. Challenges and opportunities in multichannel customer management. *J. of Serv. Research* **9**(2) 95-112.
- Ofek E., Z. Katona, M. Sarvary. 2009. “Bricks and Clicks”: The impact of product returns on the strategies of multi-channel retailers. *Marketing Science* **30**(1) 42-60.
- Pan, X., B. T. Ratchford, V. Shankar. 2002a. Can price dispersion in online markets be explained by differences in e-tailer service quality? *Journal of the Academy of Marketing Science* **30**(4) 433-445.
- Pan, X., V. Shankar, B. T. Ratchford. 2002b. Price competition between pure play vs. bricks-and-clicks e-tailers: Analytical model and empirical analysis. *Advances in Microeconomics: Economics of the Internet and e-Commerce* **11** 29-62.
- Parker, G., M. Van Alstyne. 2005. Two-Sided Network Effects: A Theory of Information Product Design. *Management Sci.* **51**(10) 1494–1504.

- Peitz, Martin, Tommaso M. Valletti. 2008. Content and advertising in the media: Pay-tv versus free-to-air. *International Journal of Industrial Organization* **26**(4) 949-965.
- Penn, M., E. K. Zalesne. "New Info Shoppers." *Wall Street Journal* 8 January 2009. Online: http://online.wsj.com/article/NA_WSJ_PUB:SB123144483005365353.html#articleTabs%3Darticle
- Prasad, A., M. Mahajan, B. Bronnenberg. 2003. Advertising versus pay-per-view in electronic media. *International Journal of Research in Marketing* **20**(1) 13-30.
- Rao, R. C. 1990. Compensating heterogeneous salesforces: some explicit solutions. *Marketing Science* **9**(4) 319-341.
- Riggins, F.J. 2003. Market segmentation and information development costs in a two-tiered fee-based and sponsorship-based Web site. *J. of Management Information Systems* **19**(3) 69-81.
- Rochet, J. C., J. Tirole. 2003. Platform competition in two-sided markets. *J. Eur. Econom. Assoc.* **1**(4) 990-1029.
- Rochet, J. C., J. Tirole. 2006. Two-sided markets: a progress report. *RAND Journal of Economics* **37**(3) 645-667.
- Shaked, A., J. Sutton. 1982. Relaxing price competition through product differentiation. *Review of Economic Studies* **49**(1) 3-13.
- Shin, J. 2005. The role of selling costs in signaling price image. *J. of Marketing Research* **42**(3) 302-312.
- Simester, D. 1995. Signaling price image using advertised prices. *Marketing Science* **14**(2) 166-188.
- Shankar, V., A. K. Smith, A. Rangaswamy. 2003. The relationship between customer satisfaction and loyalty in online and offline environments. *International J. of Research in Marketing* **20**(2) 153-175.
- Soberman, D. A. 2003. Simultaneous Signaling and Screening with Warranties. *J. of Marketing Research* **40**(5) 176-192.
- Sun, M. 2009. How does variance of product ratings matter? Working Paper, Stanford Business School.
- Sundararajan, A. 2004. Non-linear pricing of information goods. *Management Sci.* **50**(12) 1660-1673.
- Tan, T., S. Netessine. 2009. Is Tom Cruise threatened? Using Netflix prize data to examine the long tail of electronic commerce. University of Pennsylvania. Working paper.
- Teoh, S. H., C. Y. Hwang. 1991. Nondisclosure and adverse disclosure as signals of firm value. *Review of Financial Studies* **4**(2) 283-313.
- Tirole, J. 1988. *The Theory of Industrial Organization*. The MIT Press, Cambridge, MA.

- Tucker, C., J. Zhang. 2011. How does popularity information affect choices? A field experiment. *Management Science*. Forthcoming.
- Wu, S., P. Chen. 2008. Versioning and Piracy Control for Digital Information Goods. *Operations Research* **56**(1) 157-172.
- Zhu, F., X. Zhang. 2010. Impact of consumer reviews on sales: The moderating role of product and consumer characteristics. *Journal of Marketing* **74**(2) 133-148.