

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

Baojun Jiang

Olin Business School, Washington University in St. Louis, St. Louis, Missouri 63130, baojunjiang@wustl.edu

Kinshuk Jerath, Kannan Srinivasan

Tepper School of Business, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213
{kinshuk@cmu.edu, kannans@cmu.edu}

While millions of products are sold on its retail platform, Amazon.com 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. We investigate how a platform owner such as Amazon, 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).

We analyze this strategic interaction between a platform owner and an independent seller using a game-theoretic model with two types of sellers—one with high demand and one with low demand. We 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 consumer surplus in the early period, although it increases consumer surplus in the later period. We also investigate how consumer reviews influence the market outcome.

Key words: platform; retailing; long-tail products; signaling; asymmetric information; pooling equilibrium

History: Received: June 21, 2010; accepted: April 2, 2011; Eric Bradlow served as the editor-in-chief and Duncan Simester served as associate editor for this article. Published online in *Articles in Advance* July 15, 2011.

1. Introduction

Amazon.com, 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 a 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 dwarfs what is available at Walmart, the largest traditional (nonplatform) 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 best sellers, 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, and the remaining 93% are sold by independent sellers. The second column of Table 1 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.¹ 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 helps Amazon attract and retain more online customers. Because of 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

¹ For expositional ease, we will refer to the seller as “he,” the consumer as “she,” and the platform owner (Amazon) as “it” throughout the paper.

Table 1 Percentage of Products Sold by Amazon in Two Sample Product Categories

Category/subcategory	Total no. of products	% sold by Amazon	% sold by Amazon among top 100 best sellers
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
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 from Amazon.com during April 2010.

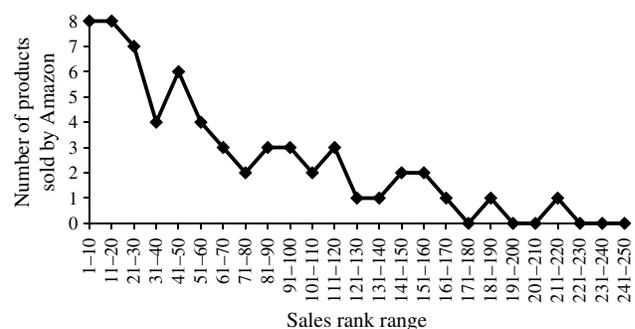
online retail platforms has become an important phenomenon, especially for long-tail products.

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 and can use these data to identify high-potential products. However, there may still be niche products for which an independent seller may have better information on demand compared with 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 best sellers 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. Although 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 best sellers.

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, although Amazon directly sells only 7% of all electronics products, it sells 64 of the top 100 best sellers. The third column of Table 1 shows that this is consistently true for other product categories. Furthermore, the percentage of products sold directly by Amazon decreases sharply as we go down the list of best sellers. Figure 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 best sellers but only 5 with sales ranks from 150 to 250.

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 is also

Figure 1 Best Sellers Sold by Amazon in the "Digital SLR Cameras" Category

not sufficiently low to ignore completely.² Interesting strategic interactions between Amazon and 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 that 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 (because 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, Amazon 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 category, where top-selling products once sold by others are now taken over by Amazon."³ Once Amazon directly procures and sells a product, it will essentially "take all of the sales away from the [independent] seller."⁴

This creates a dilemma for the high-demand seller. He may make more profits early on by selling a high volume of a product, but 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.⁵ For instance, they may devote less time and resources to dealing with consumers' inquiries about their high-demand products or related postsale services (e.g., they may offer less customization services such as gift wrapping, or they may answer product inquiries less conscientiously and with a longer time lag). These sellers may also

carry a lower inventory level and periodically create stockout situations.⁶ 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 warrant direct selling, because the observed not-so-high sales may be a result of 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 that can damage the reputation of its platform. Many anecdotes indicate that Amazon immediately terminates any sellers who 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 in deciding on how much additional (or "exceptional") service to provide beyond the standard service level. For example, gift wrapping or other customizable service options, the 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. Furthermore, the seller's promotions 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.

⁶ 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 e-mail accounts, we e-mailed two inquiries to each seller about two of their products. Two types of questions were in each e-mail: some product-specific (e.g., whether a camera's frame is metal or plastic) and some service-specific (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). Although 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.

² 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.

³ See Mikailizadeh (2010).

⁴ See McFarland (2009).

⁵ See McFarland (2010).

Whereas this 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 (because of 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 affect the answers to the above questions?

We study these 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. However, this enables the high-demand seller to mimic a low-demand seller by underinvesting 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. Interestingly, however, we find that *both* the low-demand and the high-demand sellers 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, although 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 this paper is organized as follows. In the next section, we briefly review the related literature. In §3, we develop an analytical framework to model the interaction between the platform owner and the independent seller. In §4, we first examine the complete information case; then, we analyze the incomplete information case and compare two scenarios: (1) the platform owner credibly commits to not selling the product, and (2) it retains the option to sell the product in the future. In §5, we examine the effect of consumer reviews. In §6, we discuss the robustness of our insights to alternative modeling assumptions. In §7, we conclude the paper with a short discussion.

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. Although one and occasionally more than one aspect has been studied at a time, 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 off-line consumer purchasing (Ansari et al. 2008, Bialogorsky and Naik 2003, Choi et al. 2010, Choi and Bell 2011, Neslin et al. 2006, Ofek et al. 2009, Pan et al. 2002b, Shankar et al. 2003), the impact of easier online information searches 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 through which 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 of demand 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 physical stores. 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 that sold 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 their 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, whereas 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 2011, 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. 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, although 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 because 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.⁷ 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. Although 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. In fact, our interest is in those products whose levels of demand are not already known to Amazon.

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.

⁷ 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 $\gamma^\alpha(e^{(i)})^\beta$, $\alpha > 0$, $\beta > 0$, will provide the same insights).

CONDITION C1. (i)

$$F > \frac{\theta(\gamma_H + 1/(4bk) - bc)^2 + (1 - \theta)(\gamma_L + 1/(4bk) - bc)^2}{4b} - \frac{[\theta(\gamma_H - \gamma_L) + \gamma_L + 1/(4bk) - bc]^2}{8b};$$

(ii) $F < (\gamma_H + 1/(4bk) - bc)^2/(8b)$.

CONDITION C2. $\gamma_L < \gamma_H \leq \gamma_L + 1/(2bk)$.

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 Amazon will sell the product directly if it knows $\gamma = \gamma_H$.

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 because 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 §1, 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.⁸

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. Because Amazon observes both the price and quantity sold, in essence, the seller's price-quantity pair

⁸ 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 that product.

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). Furthermore, 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.⁹ 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 App Store and Google Android Market. Note, however, that using a nonlinear 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 $q^{(1)}(p_t^{(1)}, e_t^{(1)}) = \gamma_t + e_t^{(1)} - bp_t^{(1)}$,

⁹In 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, because 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.

Table 2 Key Notations

Variables/ Symbols	Description
t	The independent seller's type; $t = "H"$ or " L ."
θ	The probability (ex ante) that the seller's type is " H ."
i	Time period; $i = 1$ 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 $p_t^{(1)}$ and $q_t^{(1)}$.
$\Pi_A^{(i)}$	Amazon's expected profit in period i .
Π_A	Amazon's overall expected profit for both periods $i = 1$ or 2 .
$\Pi_t^{(i)}$	Type t seller's profit in period i .
Π_t	Type t seller's overall profit for both periods $i = 1$ or 2 .
$\hat{\cdot}$	The hat over a variable indicates the case of complete information.
$\bar{\cdot}$	The bar over a variable indicates the case of no entry threat by Amazon.
sep, pool	These subscripts indicate the separating and pooling outcomes, respectively.
rev	This subscript indicates that the variable is for the case with seller reviews.

and both the seller and Amazon realize their respective profits.

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 2.

4. Analysis

We organize our analysis in the following way. In §4.1, we analyze a benchmark case of complete information, in which Amazon knows the true demand (i.e., the seller's type). In §4.2, we analyze another benchmark case of asymmetric information and unobserv-

able service, in which Amazon contractually commits to not selling the product. In §4.3, we examine our focal case of asymmetric information and unobservable service in which Amazon retains the option to sell the product.

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 the product 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

$$\hat{\Pi}_L = \sum_{i=1}^2 \{(\gamma_L + \hat{e}_L^{(i)} - b\hat{p}_L^{(i)})[\hat{p}_L^{(i)} - c - \hat{f} - k(\hat{e}_L^{(i)})^2]\},$$

where $c > 0$ is the product's marginal cost.¹⁰ Solving the first-order conditions, one easily finds the L -type seller's optimal (first-best) service level and price: $\hat{e}_L^{*(i)} = 1/(2bk)$ and $\hat{p}_L^{*(i)} = (\gamma_L + b(c + \hat{f}) + 3/(4bk))/(2b)$, respectively. The corresponding profit is given by

$$\hat{\Pi}_L^*(\hat{f}) = \frac{(\gamma_L - b(c + \hat{f})) + 1/(4bk)^2}{2b}.$$

Amazon's profit in the case of an L -type seller is given by $\hat{\Pi}_{A,L}(\hat{f}) = \sum_{i=1}^2 \{\hat{f}\hat{q}_L^{(i)}(\hat{f})\}$, 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 $\hat{f}^* = (\gamma_L - bc + 1/(4bk))/(2b)$, and its profit from the L -type seller is $\hat{\Pi}_{A,L}^* = (\gamma_L - bc + 1/(4bk))^2/(4b)$.

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 $e_A^{*(i)} = 1/(2bk)$ and $p_A^{*(i)} = (\gamma_H + bc + 3/(4bk))/(2b)$, respectively, with a corresponding total profit of $\hat{\Pi}_{A,H}^* = 2[(\gamma_H - bc + 1/(4bk))^2/(4b) - F]$.

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.¹¹ 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 $\bar{\Pi}_t = \sum_{i=1}^2 \{(\gamma_t + \bar{e}_t^{(i)} - b\bar{p}_t^{(i)})[\bar{p}_t^{(i)} - c - \bar{f} - k(\bar{e}_t^{(i)})^2]\}$, 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 this paper, 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.¹² The equilibrium outcome (service levels, prices, fees, and overall profits for a seller of type t and for Amazon) are provided in the appendix.

4.3. The Platform Owner Retains Its Entry Option

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).

4.3.1. 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 $f < (\gamma_L - bc + 1/(4bk))/b$. This is because at such a fee, both types of sellers sell on the platform, and

¹⁰ For analytical simplicity and without loss of generality, we assume the discount rate to be 1 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.

¹¹ Amazon can make a legally binding credible commitment of this nature by putting in its seller agreement verbiage such as "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."

¹² All proofs in this paper are relegated to the appendix.

the H -type strictly prefers to mimic 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 $(\gamma_L - bc + 1/(4bk))/b \leq f < (\gamma_H - bc + 1/(4bk))/b$ will induce a separating equilibrium because with such a fee, only the H -type seller will profitably sell a positive quantity (the L -type will not enter the market).¹³ 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 during which Amazon will sell the product directly.¹⁴ 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 compute f_{sep}^* , $\Pi_{t, sep}^*$, and $\Pi_{A, sep}^*$. In §4.3.3, we specify in Proposition 1 the condition under which this separating equilibrium is realized.

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

We now consider the case of $f < (\gamma_L - bc + 1/(4bk))/b$, in which both types of sellers will sell on Amazon's platform. Note that if the realized equilibrium corresponds to $f < (\gamma_L - bc + 1/(4bk))/b$, then a pooling equilibrium arises because given such a fee, the H -type seller strictly prefers to mimic the L -type to avoid Amazon's entry in the second period. Because 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 electronic companion (available

as part of the online version that can be found at <http://mktsci.pubs.informs.org/>), we show that such refinements lead to the following unique pooling equilibrium: in the first period, both types of sellers choose $p_t^{*(1)}(f) = (\gamma_L + b(c + f) + 3/(4bk))/(2b)$ and sell a quantity of $q_t^{*(1)}(f) = (\gamma_L - b(c + f) + 1/(4bk))/2$. 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 $e_t^{*(2)} = 1/(2bk)$ and $p_t^{*(2)}(f) = (\gamma_t + b(c + f) + 3/(4bk))/(2b)$, which are the same as (3) and (5) in the appendix, respectively. The seller's second-period profit is given by $\Pi_t^{*(2)}(f) = [(\gamma_t - b(c + f) + 1/(4bk))]^2/(4b)$.

Amazon observes the seller's first-period price $p_t^{(1)}$ and unit sales $q_t^{(1)}$ because it processes all orders on its platform. In effect, $p_t^{(1)}$ 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 because 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 Amazon, knowing that he is an L -type, 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 because 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

¹³ The upper bound is needed because with $f \geq (\gamma_H - bc + 1/(4bk))/b$, neither type of seller can profitably sell on Amazon.

¹⁴ 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 $p_L < (\gamma_L + e_L)/b$. If $f \geq (\gamma_L - bc + 1/(4bk))/b$, the L -type's profit margin becomes

$$\begin{aligned} p_L - c - f - ke_L^2 &< \frac{\gamma_L + e_L}{b} - c - \frac{\gamma_L - bc + 1/(4bk)}{b} - ke_L^2 \\ &= \frac{e_L - bke_L^2 + 1/(4bk)}{b} = -k \left(e_L - \frac{1}{2bk} \right)^2 < 0. \end{aligned}$$

Hence, he will not sell on Amazon if its fee is so high.

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. Because 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 §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.

4.3.3. 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 2 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, and 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 are likely to 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, and 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 because 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 paper is the pooling equilibrium.

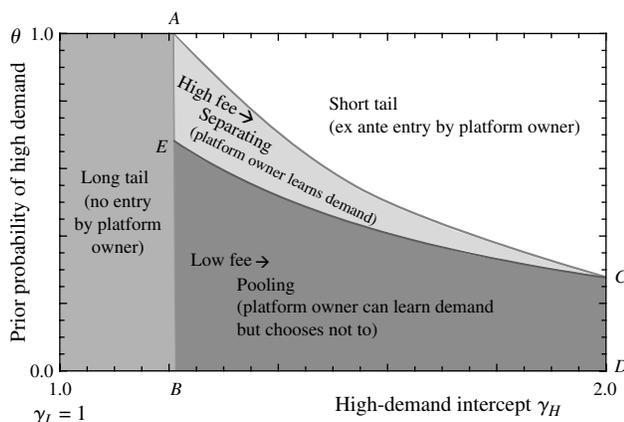
Now, we examine how Amazon's entry option affects its own profit.

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

Figure 2 Market Outcomes in (γ_H, θ) Parameter Space



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, and 15% for toys and video games. 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 account 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 will set 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 because 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 as a result of 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. If γ_H is very high, however, 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. Interestingly, however, 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 as a result of 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. However, 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.

5. Effect of Consumer Reviews

Although Amazon may not directly observe the seller's service level, it can solicit reviews from the seller's first-period customers (because 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.¹⁵

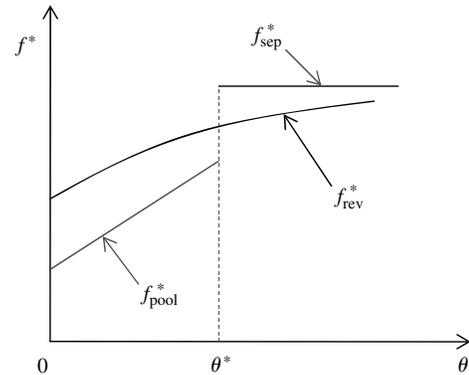
Here, we assume the extreme case that consumer reviews fully reveal the seller's service level. After the first period, Amazon will acquire such reviews to determine the seller's service level and hence correctly infer the seller's type. Note that the fully revealing reviews do not imply a full (complete)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

Because 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 $e_i^* = 1/(2bk)$ and $p_i^*(f_{\text{rev}}) = (\gamma_i + b(c + f_{\text{rev}}) + 3/(4bk))/(2b)$, which yields a first-period profit of $\Pi_i^{*(1)} = [\gamma - b(c + f_{\text{rev}}) + 1/(4bk)]^2/(4b)$.

In the case of an H -type seller, Amazon's first-period profit is $f_{\text{rev}} \cdot (\gamma_H - b(c + f_{\text{rev}}) + 1/(4bk))/2$. In the second period, Amazon will sell the product directly, making a maximum profit of $(\gamma_H - bc + 1/(4bk))^2/(4b) - F$ at the optimal price of $p_A^* = (\gamma_H + bc + 3/(4bk))/(2b)$. With the L -type seller, Amazon makes the same profit of $f_{\text{rev}} \cdot$

¹⁵ 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 learn the seller's type after the first period.

Figure 3 Effect of Reviews on Amazon's Fee



$(\gamma_L - b(c + f_{\text{rev}}) + 1/(4bk))/2$ for each period. Therefore, Amazon's expected total profit for both periods is given by¹⁶

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

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

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

and

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

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).*

Proposition 6, illustrated in Figure 3, 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 because the H -type

¹⁶ We have assumed that $\theta < (2(\gamma_L - bc + 1/(4bk)))/(\gamma_H - bc + 1/(4bk))$; 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.

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 because 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. However, 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, which is given by $\Pi_{A,rev}^* - \Pi_A^*$. 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 .

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 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 trade-offs between the platform owner and the seller, and hence our key results remain qualitatively the same (although 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 (which has three components: the wholesale price, the fee paid to Amazon, and the service cost) to the seller. We find through numerical examples that such intuition is indeed correct. Because 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 and 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 smaller than 1) in the second period. Although 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 will exit the market to encourage entry by the third-party sellers.

Fourth, although 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. In essence, 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 nonlinear 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 electronic companion, we show that our focal pooling equilibrium still exists in certain parameter regions even when a general nonlinear 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 will 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 thus 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 because Amazon has the same information about the seller's type at the beginning of the first and second periods (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 who would otherwise not enter the platform.

7. Conclusions

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 hard-to-find 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 that a high-demand seller has an 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. Because of 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, although 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 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). However, even good consumer reviews are unlikely to fully reveal the seller's service level, because such reviews tend to reveal postsale service levels rather than presale services. If there is mediocre presale 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 nonplatform 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 it knows the product 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 it 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 type H or L) 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). 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.

8. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://mktsci.pubs.informs.org/>.

Acknowledgments

This paper is based on the second essay of B. Jiang's doctoral dissertation at Carnegie Mellon University. The authors are grateful to Nitin Mehta for extensive discussions that have tremendously improved this paper. The authors also thank seminar participants at the Rotman School of Management at the University of Toronto, the Ross School of Business at the University of Michigan, the Olin Business School at Washington University in St. Louis, INSEAD, Tsinghua University, Peking University, and the editor, area editor, and two anonymous reviewers.

Appendix

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 because 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 (FOCs) are

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

$$\frac{\partial \bar{\Pi}_t^{(i)}}{\partial \bar{e}_t^{(i)}} = (\gamma_t + \bar{e}_t^{(i)} - b\bar{p}_t^{(i)}) \cdot (-2k\bar{e}_t^{(i)}) + [\bar{p}_t^{(i)} - c - \bar{f} - k(\bar{e}_t^{(i)})^2] = 0. \quad (3)$$

From (2), we obtain $\bar{p}_t^{(i)} = (\gamma_t + \bar{e}_t^{(i)} + b(c + \bar{f} + k(\bar{e}_t^{(i)})^2))/(2b)$. Substituting this into (3), we then solve for $\bar{e}_t^{(i)}$:

$$\frac{\partial \bar{\Pi}_t^{(i)}}{\partial \bar{e}_t^{(i)}} = k \left(\bar{e}_t^{(i)} - \frac{1}{2bk} \right) \{ bk(\bar{e}_t^{(i)})^2 - \bar{e}_t^{(i)} - [\gamma_t - b(c + \bar{f})] \} = 0.$$

Thus, the potential FOC solutions are

$$\bar{e}_t^{*(i)} = \frac{1}{2bk}, \quad (4)$$

$$\bar{e}_t^{*(i)} = \frac{1 \pm \sqrt{1 + 4bk[\gamma_t - b(c + \bar{f})]}}{2bk}. \quad (5)$$

We eliminate (5) since simple algebra shows that 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 (4) is the local maximum and (5) is a saddle point. It is easy to show that the boundary of $\bar{e}_t^{(i)} = 0$ and $\bar{p}_t^{(i)} = 0$ yields a lower profit than (4). Hence, (4) is the seller's (global) profit-maximizing service level, with the corresponding optimal price and profit respectively given by (6) and (7) below:

$$\bar{p}_t^{*(i)}(\bar{f}) = \frac{\gamma_t + b(c + \bar{f}) + 3/(4bk)}{2b}, \quad (6)$$

$$\bar{\Pi}_t^*(\bar{f}) = \bar{\Pi}_t^{*(1)} + \bar{\Pi}_t^{*(2)} = \frac{[\gamma_t - b(c + \bar{f}) + 1/(4bk)]^2}{2b}. \quad (7)$$

Lemma 1 immediately follows. \square

Equilibrium Outcome When Amazon Commits to No Entry (§4.2)

The optimal service levels, prices, and overall profit for a seller of type t are given by (4), (6), and (7), respectively. Amazon's expected total profit is given by $\bar{\Pi}_A(\bar{f}) = \sum_{i=1}^2 \{\bar{f}[\theta \bar{q}_H^{(i)}(\bar{f}) + (1-\theta)\bar{q}_L^{(i)}(\bar{f})]\}$, 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 (4) and (6) into $\bar{q}_t^{(i)}(\bar{f})$ leads to $\bar{\Pi}_A(\bar{f}) = \bar{f}[\theta\gamma_H + (1-\theta)\gamma_L - b(c + \bar{f}) + 1/(4bk)]$. Amazon's equilibrium fee and profit are¹⁷

$$\bar{f}^* = \frac{\theta\gamma_H + (1-\theta)\gamma_L - bc + 1/(4bk)}{2b}, \quad (8)$$

$$\bar{\Pi}_A^* = \frac{[\theta\gamma_H + (1-\theta)\gamma_L - bc + 1/(4bk)]^2}{4b}. \quad (9)$$

Substituting (8) into (6) and (7), we obtain the equilibrium outcome for type $t \in \{L, H\}$ seller:

$$\bar{p}_t^{*(1)} = \bar{p}_t^{*(2)} = \frac{2\gamma_t + \theta\gamma_H + (1-\theta)\gamma_L + bc + 7/(4bk)}{4b}, \quad (10)$$

$$\bar{\Pi}_t^* = \frac{[2\gamma_t - \theta\gamma_H - (1-\theta)\gamma_L - bc + 1/(4bk)]^2}{8b}. \quad (11)$$

Separating Outcome

Amazon's expected profit is expressed as

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

A proof very similar to that for Lemma 1 shows that $e_A^* = 1/(2bk)$ and $p_A^* = (\gamma_H + bc + 3/(4bk))/(2b)$ (because of space considerations, we exclude it from this paper). With this, we can rewrite Amazon's expected profit as a function of only f_{sep} , where, as discussed before, $(\gamma_L - bc + 1/(4bk))/b \leq f_{\text{sep}} < (\gamma_H - bc + 1/(4bk))/b$:

$$\Pi_{A, \text{sep}}(f_{\text{sep}}) = \theta f_{\text{sep}} \cdot \frac{\gamma_H - b(c + f_{\text{sep}}) + 1/(4bk)}{2} + \theta \left[\frac{(\gamma_H - bc + 1/(4bk))^2}{4b} - F \right]. \quad (12)$$

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, because doing so yields a profit of $(\gamma_H - bc + 1/(4bk))^2/(4b) - F$ rather than the maximum potential profit of $(\gamma_H - bc + 1/(4bk))^2/(8b)$ 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 $f < (\gamma_L - bc + 1/(4bk))/b$.

¹⁷ Here, we have implicitly assumed a nonboundary 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 $f^* = (\gamma_H - bc + 1/(4bk))/(2b)$). One can easily show that such a boundary solution requires both

$$\theta > \left(\frac{\gamma_L - bc + 1/(4bk)}{\gamma_H - \gamma_L} \right)^2 \quad \text{and} \quad \gamma_H > 2\gamma_L - bc + \frac{1}{4bk}.$$

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 (12), if a separating equilibrium is realized, Amazon's fee must be given by (13), its expected profit by (14), and the seller's profit by (15) and (16). The two forms in (13)–(15) 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.

$$f_{\text{sep}}^* = \begin{cases} \frac{\gamma_L - bc + 1/(4bk)}{b}, & \text{if } \gamma_H < 2\gamma_L - bc + \frac{1}{4bk}, \\ \frac{\gamma_H - bc + 1/(4bk)}{2b}, & \text{if } \gamma_H \geq 2\gamma_L - bc + \frac{1}{4bk}; \end{cases} \quad (13)$$

$$\Pi_{A, \text{sep}}^* = \begin{cases} \theta \left[\frac{(\gamma_L - bc + 1/(4bk))(\gamma_H - \gamma_L)}{2b} + \frac{(\gamma_H - bc + 1/(4bk))^2}{4b} - F \right], & \text{if } \gamma_H < 2\gamma_L - bc + \frac{1}{4bk} \\ \theta \left[\frac{3(\gamma_H - bc + 1/(4bk))^2}{8b} - F \right], & \text{if } \gamma_H \geq 2\gamma_L - bc + \frac{1}{4bk}; \end{cases} \quad (14)$$

$$\Pi_{H, \text{sep}}^* = \begin{cases} \frac{(\gamma_H - \gamma_L)^2}{4b}, & \text{if } \gamma_H < 2\gamma_L - bc + \frac{1}{4bk}, \\ \frac{(\gamma_H - bc + 1/(4bk))^2}{16b}, & \text{if } \gamma_H \geq 2\gamma_L - bc + \frac{1}{4bk}; \end{cases} \quad (15)$$

$$\Pi_{L, \text{sep}}^* = 0. \quad (16)$$

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 (17) and (18) for the L -type seller and (19)–(22) for the H -type seller:

$$e_{L, \text{pool}}^{*(1)} = e_{L, \text{pool}}^{*(2)} = \frac{1}{2bk}, \quad (17)$$

$$p_{L, \text{pool}}^{*(1)} = p_{L, \text{pool}}^{*(2)} = \frac{\gamma_L + b(c + f) + 3/(4bk)}{2b}, \quad (18)$$

$$e_{H, \text{pool}}^{*(1)} = \gamma_L + \frac{1}{2bk} - \gamma_H, \quad (19)$$

$$p_{H, \text{pool}}^{*(1)} = \frac{\gamma_L + b(c + f) + 3/(4bk)}{2b}, \quad (20)$$

$$e_{H, \text{pool}}^{*(2)} = \frac{1}{2bk}, \quad (21)$$

$$p_{H, \text{pool}}^{*(2)} = \frac{\gamma_H + b(c + f) + 3/(4bk)}{2b}. \quad (22)$$

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

$$\Pi_{L, \text{pool}}^*(f) = \frac{[\gamma_L - b(c + f) + 1/(4bk)]^2}{2b}, \quad (23)$$

$$\begin{aligned} \Pi_{H, \text{pool}}^*(f) &= \Pi_{H, \text{pool}}^{*(1)}(f) + \Pi_{H, \text{pool}}^{*(2)}(f) \\ &= [p_{H, \text{pool}}^{*(1)} - c - f - k(e_{H, \text{pool}}^{*(1)})^2](\gamma_H + e_{H, \text{pool}}^{*(1)} - bp_{H, \text{pool}}^{*(1)}) \\ &\quad + \frac{[\gamma_H - b(c + f) + 1/(4bk)]^2}{4b}. \end{aligned} \quad (24)$$

Amazon's expected total profit is given by $\Pi_{A,\text{pool}}(f) = \Pi_{A,\text{pool}}^{(1)}(f) + \Pi_{A,\text{pool}}^{(2)}(f)$. It chooses f to maximize its expected profit. Amazon's optimal pooling fee and profit are given by¹⁸

$$f_{\text{pool}}^* = \frac{\theta(\gamma_H - \gamma_L)/2 + \gamma_L - bc + 1/(4bk)}{2b}, \quad (25)$$

$$\Pi_{A,\text{pool}}^* = \frac{[\theta(\gamma_H - \gamma_L)/2 + \gamma_L - bc + 1/(4bk)]^2}{4b}. \quad (26)$$

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 $\Pi_{A,\text{sep}}^* \geq \Pi_{A,\text{pool}}^*$ and a pooling equilibrium if $\Pi_{A,\text{sep}}^* < \Pi_{A,\text{pool}}^*$.

Case 1. $\gamma_H \geq 2\gamma_L - bc + 1/(4bk)$.

We examine two subcases. First, we consider the condition of

$$\theta > \frac{2(\gamma_L - bc + 1/(4bk))}{\gamma_H - \gamma_L}.$$

As discussed before, a pooling equilibrium requires $f < (\gamma_L - bc + 1/(4bk))/b$ such that the L -type will profitably sell a positive quantity. However, if

$$\theta > \frac{2(\gamma_L - bc + 1/(4bk))}{\gamma_H - \gamma_L},$$

the FOC solution

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

implies that Amazon's best pooling outcome corresponds to a fee $f \rightarrow (\gamma_L - bc + 1/(4bk))/b$. As $\rightarrow (\gamma_L - bc + 1/(4bk))/b$,

$$\Pi_{A,\text{pool}}(f) \rightarrow \frac{\theta(\gamma_H - \gamma_L)[\gamma_L - bc + 1/(4bk)]}{2b}.$$

From (14) and assumption C1(ii), we get

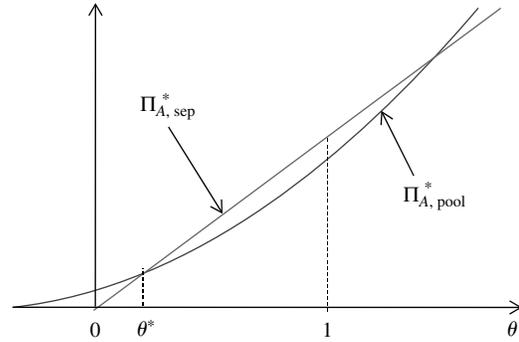
$$\begin{aligned} \Pi_{A,\text{sep}}^* &= \theta \left[\frac{3(\gamma_H - bc + 1/(4bk))^2}{8b} - F \right] > \frac{\theta(\gamma_H - bc + 1/(4bk))^2}{4b} \\ &> \frac{\theta(\gamma_H - \gamma_L)[\gamma_L - bc + 1/(4bk)]}{2b}. \end{aligned}$$

Note that the last inequality is proved by expanding the terms:

$$\begin{aligned} \frac{\theta(\gamma_H - bc + 1/(4bk))^2}{4b} &> \frac{\theta(\gamma_H - \gamma_L)[\gamma_L - bc + 1/(4bk)]}{2b} \\ \Leftrightarrow (\gamma_H - \gamma_L)^2 - (\gamma_L - bc + 1/(4bk))^2 &> 0 \\ \Leftrightarrow \gamma_H - \gamma_L > \gamma_L - bc + \frac{1}{4bk}, \end{aligned}$$

¹⁸ In the proof of Proposition 1, we show that Amazon's pooling profit at the boundary $f \rightarrow (\gamma_L - bc + 1/(4bk))/b$ is always lower than its separating profit. Hence, (27) is the only possible pooling equilibrium fee.

Figure A.1 Amazon's Pooling and Separating Profits



which is true under Case 1. Thus, we conclude that the separating equilibrium is realized if

$$\theta > \frac{2(\gamma_L - bc + 1/(4bk))}{\gamma_H - \gamma_L}.$$

Second, we consider the condition of $\theta \leq 2(\gamma_L - bc + 1/(4bk))/(\gamma_H - \gamma_L)$, which means that the best pooling outcome occurs at the FOC point (25). We compare Amazon's profits from the potential separating equilibrium (14) and the potential pooling equilibrium (26). The separating equilibrium is realized if and only if $\Pi_{A,\text{sep}}^* \geq \Pi_{A,\text{pool}}^*$ or

$$\begin{aligned} \theta \left[\frac{3(\gamma_H - bc + 1/(4bk))^2}{8b} - F \right] \\ - \frac{[\theta(\gamma_H - \gamma_L)/2 + \gamma_L - bc + 1/(4bk)]^2}{4b} \geq 0. \end{aligned}$$

We plot both $\Pi_{A,\text{sep}}^*$ and $\Pi_{A,\text{pool}}^*$ 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

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

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

$$\begin{aligned} \theta \left[\frac{3(\gamma_H - bc + 1/(4bk))^2}{8b} - F \right] \\ = \frac{[\theta(\gamma_H - \gamma_L)/2 + \gamma_L - bc + 1/(4bk)]^2}{4b}, \end{aligned}$$

is given by

$$\begin{aligned} \theta^* &= (3(\gamma_H - bc + 1/(4bk))^2 - 8bF - 2(\gamma_H - \gamma_L) \\ &\quad \cdot (\gamma_L - bc + 1/(4bk)) - ([3(\gamma_H - bc + 1/(4bk))^2 - 8bF \\ &\quad - 2(\gamma_H - \gamma_L)(\gamma_L - bc + 1/(4bk))]^2 - 4(\gamma_H - \gamma_L)^2 \\ &\quad \cdot (\gamma_H - bc + 1/(4bk))^2)^{1/2} \cdot (\gamma_H - \gamma_L)^{-2}. \end{aligned}$$

Case 2. $\gamma_H < 2\gamma_L - bc + 1/(4bk)$.

The separating equilibrium is realized if and only if $\Pi_{A,sep}^* \geq \Pi_{A,pool}^*$ or

$$\begin{aligned} \theta \left[\frac{(\gamma_L + 1/(4bk) - bc)(\gamma_H - \gamma_L)}{2b} + \frac{(\gamma_H - bc + 1/(4bk))^2}{4b} - F \right] \\ - \frac{[\theta(\gamma_H - \gamma_L)/2 + \gamma_L - bc + 1/(4bk)]^2}{4b} \geq 0. \end{aligned}$$

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

$$\begin{aligned} \theta^* &= (2(\gamma_H - bc + 1/(4bk))^2 - 8bF + 2(\gamma_H - \gamma_L) \\ &\quad \cdot (\gamma_L - bc + 1/(4bk)) - ([2(\gamma_H - bc + 1/(4bk))^2 - 8bF \\ &\quad + 2(\gamma_H - \gamma_L)(\gamma_L - bc + 1/(4bk))]^2 - 4(\gamma_H - \gamma_L)^2 \\ &\quad \cdot (\gamma_H - bc + 1/(4bk))^2)^{1/2} \cdot (\gamma_H - \gamma_L)^{-2}. \quad \square \end{aligned}$$

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 (8) and (25), it is simple to show $f_{pool}^* < \bar{f}^*$. Comparing Amazon's pooling equilibrium profit given by (26) with its no-threat-of-entry profit given by (9), one easily shows that $\Pi_{A,pool}^* < \bar{\Pi}_A^*$. \square

PROOF OF PROPOSITION 3. Please refer to the electronic companion for the 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 (4) and (6) 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

$$\overline{CS}_t^{(1)}(\bar{f}^*) = \overline{CS}_t^{(2)}(\bar{f}^*) = \frac{[\gamma_t - b(c + \bar{f}^*) + 1/(4bk)]^2}{8b}.$$

When Amazon keeps its entry option, the pooling equilibrium is realized if $\theta < \theta^*$; the second-period consumer surplus is

$$CS_t^{(2)}(f_{pool}^*) = \frac{[\gamma_t - b(c + f_{pool}^*) + 1/(4bk)]^2}{8b}.$$

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 because 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, because the H -type exactly mimics the L -type. Thus,

$$CS_H^{(1)}(f_{pool}^*) = CS_L^{(1)}(f_{pool}^*) = \frac{[\gamma_L - b(c + f_{pool}^*) + 1/(4bk)]^2}{8b}.$$

Using Equations (8) and (25), one can then easily show that $CS_H^{(1)}(f_{pool}^*) < \overline{CS}_H^{(1)}(\bar{f}^*)$. \square

PROOF OF PROPOSITION 5. This follows immediately from (25) and (13).

PROOF OF PROPOSITION 6.

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

If $\gamma_H < 2\gamma_L - bc - 1/4bk$,

$$\begin{aligned} f_{rev}^* - f_{sep}^* &= \frac{\theta(\gamma_H - bc + 1/(4bk)) + 2(1 - \theta)(\gamma_L - bc + 1/(4bk))}{2b(2 - \theta)} \\ &\quad - \frac{\gamma_L - bc + 1/(4bk)}{b} \\ &= \frac{1}{2b(2 - \theta)} \times [(\theta(\gamma_H - bc + 1/(4bk)) \\ &\quad + 2(1 - \theta)(\gamma_L - bc + 1/(4bk)) \\ &\quad - 2(2 - \theta)(\gamma_L - bc + 1/(4bk)))] \\ &= \frac{\theta(\gamma_H - bc + 1/(4bk)) - 2(\gamma_L - bc + 1/(4bk))}{2b(2 - \theta)} < 0. \end{aligned}$$

If $\gamma_H \geq 2\gamma_L - bc - 1/4bk$,

$$\begin{aligned} f_{rev}^* - f_{sep}^* &= \frac{\theta(\gamma_H - bc + 1/(4bk)) + 2(1 - \theta)(\gamma_L - bc + 1/(4bk))}{2b(2 - \theta)} \\ &\quad - \frac{\gamma_H - bc + 1/(4bk)}{2b} \\ &= \frac{1}{2b(2 - \theta)} \times [(\theta(\gamma_H - bc + 1/(4bk)) \\ &\quad + 2(1 - \theta)(\gamma_L - bc + 1/(4bk)) \\ &\quad - (2 - \theta)(\gamma_H - bc + 1/(4bk)))] \\ &= \frac{(1 - \theta)(\gamma_L - \gamma_H)}{b(2 - \theta)} < 0. \end{aligned}$$

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

References

- Anderson, C. 2006. *The Long Tail: Why the Future of Business Is Selling Less of More*. Hyperion, New York.
- Ansari, A., C. F. Mela, S. A. Neslin. 2008. Customer channel migration. *J. Marketing Res.* 45(1) 60–76.
- Araujo, A., D. Gottlieb, H. Moreira. 2008. A model of mixed signals with applications to countersignaling and the GEC. *RAND J. Econom.* 38(4) 1020–1043.
- Armstrong, M. 2006. Competition in two-sided markets. *RAND J. Econom.* 37(3) 668–691.
- Bakos, J. Y. 1997. Reducing buyer search costs: Implications for electronic marketplaces. *Management Sci.* 43(12) 1676–1692.

- Baye, M. R., J. R. J. Gatti, P. Kattuman, J. Morgan. 2007. A dashboard for online pricing. *Calif. Management Rev.* 50(1) 202–216.
- Biyalogorsky, E., P. Naik. 2003. Clicks and mortar: The effect of online activities on offline sales. *Marketing Lett.* 14(1) 21–32.
- Brynjolfsson, E., M. D. Smith. 2000. Frictionless commerce? A comparison of internet and conventional retailers. *Management Sci.* 46(4) 563–585.
- 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 Sci.* 49(11) 1580–1596.
- Brynjolfsson, E., Y. Hu, M. D. Smith. 2006. From niches to riches: The anatomy of the long tail. *Sloan Management Rev.* 47(4) 67–71.
- Chen, Y., J. Xie. 2005. Third-party product review and firm marketing strategy. *Marketing Sci.* 24(2) 218–240.
- Cho, I. K., D. M. Kreps. 1987. Signaling games and stable equilibria. *Quart. J. Econom.* 102(2) 179–222.
- Choi, J., D. R. Bell. 2011. Preference minorities and the Internet. *J. Marketing Res.* 48(4) 670–682.
- Choi, J., S. K. Hui, D. Bell. 2010. Spatio-temporal analysis of imitation behavior across new buyers at an online grocery retailer. *J. Marketing Res.* 47(1) 65–79.
- Coughlan, A. T., B. Wernerfelt. 1989. On credible delegation by oligopolists: A discussion of distribution channel management. *Management Sci.* 35(2) 226–239.
- Desai, P. S. 2000. Multiple messages to retain retailers: Signaling new product demand. *Marketing Sci.* 19(4) 381–389.
- Desai, P. S., K. Srinivasan. 1995. Demand signaling under unobservable effort in franchising: Linear and non-linear price contracts. *Management Sci.* 41(10) 1608–1623.
- Desai, P., O. Koenigsberg, D. Purohit. 2004. Strategic decentralization and channel coordination. *Quant. Marketing Econom.* 2(1) 5–22.
- Desiraju, R., S. Moorthy. 1997. Managing a distribution channel under asymmetric information with performance requirements. *Management Sci.* 43(12) 1628–1644.
- Elberse, A. 2008. Should you invest in the long tail? *Harvard Bus. Rev.* 86(7/8) 88–96.
- Eisenmann, T., G. Parker, M. W. Van Alstyne. 2006. Strategies for two-sided markets. *Harvard Bus. Rev.* 84(October) 92–101.
- Feltovich, N., R. Harbaugh, T. To. 2002. Too cool for school? Signalling and countersignalling. *RAND J. Econom.* 33(4) 630–649.
- Holmstrom, B. 1979. Moral hazard and observability. *Bell J. Econom.* 10(1) 74–91.
- Iyer, G., M. Villas-Boas. 2003. A bargaining theory of distribution channels. *J. Marketing Res.* 40(1) 80–100.
- Jerath, K., Z. J. Zhang. 2010. Store within a store. *J. Marketing Res.* 47(4) 748–763.
- Jeuland, A. P., S. M. Shugan. 1983. Managing channel profits. *Marketing Sci.* 2(3) 239–272.
- Jiang, B., K. Srinivasan. 2011. Impact of online consumer reviews on competitive marketing strategies. Working paper, Carnegie Mellon University, Pittsburgh.
- Kuksov, D., Y. Xie. 2010. Pricing, frills, and customer ratings. *Marketing Sci.* 29(5) 925–943.
- Lal, R., R. Staelin. 1986. Salesforce compensation plans in environments with asymmetric information. *Marketing Sci.* 5(3) 179–198.
- Mayzlin, D., J. Shin. 2011. Uninformative advertising as an invitation to search. *Marketing Sci.* 30(4) 666–685.
- McFarland, A. 2009. Amazon lets us pay them to grow. *Adam McFarland* (blog), October 18, <http://www.adam-mcfarland.net/2009/10/18/amazon-lets-us-pay-them-to-grow/>.
- McFarland, A. 2010. How Amazon exploits the "mid-tail." *Adam McFarland* (blog), July 6, <http://www.adam-mcfarland.net/2010/07/06/how-amazon-exploits-the-mid-tai/>.
- McGuire, T. W., R. Staelin. 1983. An industry equilibrium analysis of downstream vertical integration. *Marketing Sci.* 2(2) 161–191.
- Mikailizadeh, M. 2010. Amazon's world domination plan. *Motley Fool* (February 9), <http://www.fool.com/investing/general/2010/02/09/amazons-world-domination-plan.aspx>.
- Moorthy, K. S. 1988. Strategic decentralization in channels. *Marketing Sci.* 7(4) 335–355.
- Moorthy, S., K. Srinivasan. 1995. Signaling quality with a money-back guarantee: The role of transaction costs. *Marketing Sci.* 14(4) 442–466.
- 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. Service Res.* 9(2) 95–112.
- Ofek, E., Z. Katona, M. Sarvary. 2009. "Bricks and clicks": The impact of product returns on the strategies of multichannel retailers. *Marketing Sci.* 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? *J. Acad. Marketing Sci.* 30(4) 433–445.
- Pan, X., V. Shankar, B. T. Ratchford. 2002b. Price competition between pure play versus bricks-and-clicks e-tailers: Analytical model and empirical analysis. M. R. Baye, ed. *The Economics of the Internet and E-Commerce*. Advances in Applied Microeconomics, Vol. 11. Elsevier, Oxford, UK, 29–62.
- Parker, G., M. W. Van Alstyne. 2005. Two-sided network effects: A theory of information product design. *Management Sci.* 51(4) 1494–1504.
- Rao, R. C. 1990. Compensating heterogeneous salesforces: Some explicit solutions. *Marketing Sci.* 9(4) 319–341.
- Rochet, J. C., J. Tirole. 2003. Platform competition in two-sided markets. *J. Eur. Econom. Assoc.* 1(4) 990–1029.
- Shin, J. 2005. The role of selling costs in signaling price image. *J. Marketing Res.* 42(3) 302–312.
- Simester, D. 1995. Signalling price image using advertised prices. *Marketing Sci.* 14(2) 166–188.
- Shankar, V., A. K. Smith, A. Rangaswamy. 2003. Customer satisfaction and loyalty in online and offline environments. *Internat. J. Res. Marketing* 20(2) 153–175.
- Soberman, D. A. 2003. Simultaneous signaling and screening with warranties. *J. Marketing Res.* 40(2) 176–192.
- Tan, T., S. Netessine. 2009. Is Tom Cruise threatened? Using Netflix prize data to examine the long tail of electronic commerce. Working paper, University of Pennsylvania, Philadelphia.
- Teoh, S. H., C. Y. Hwang. 1991. Nondisclosure and adverse disclosure as signals of firm value. *Rev. Financial Stud.* 4(2) 283–313.
- Tucker, C., J. Zhang. 2011. How does popularity information affect choices? A field experiment. *Management Sci.* 57(5) 828–842.