Amazon and the Evolution of Retail

Tommaso Bondi
Cornell University

Luís Cabral
New York University and *CEPR*

June 2022

Abstract. The growth of Amazon and other online retailers questions the survival of bricks-and-mortar retail. We show that, in response to the online trend, offline retailers — especially smaller ones — optimally follow a specialization strategy, in particular specialization in narrow niches. This may lead to offline markets being more niche-concentrated than online ones, contrary to the conventional wisdom of the "embarrassment of niches" induced by online sales. The intuition for this result is that the growth of online platforms like Amazon hurts all bricks-and-mortar stores, but it especially hurts large stores selling popular-appeal items. We discuss this and other relevant comparative statics based on a simple model of consumer demand and retail design. We develop various extensions, including pricing, consumer eclecticism, offline amenities, and the role of offline-to-offline competition. We also show theoretically that offline-store specialization benefits consumers, and that in equilibrium bricks-and-mortar stores fall short of what consumers would prefer in terms of specialization. Finally, we complement our theoretical analysis with empirical evidence building on a large proprietary dataset obtained from a major US publisher detailing all sales to book retailers (both online and offline) over the 2016-2019 period. The evidence is consistent with the prediction that independent bookstores, and especially smaller ones, are more likely to sell niche-genre titles, to the point that offline sales are less top-heavy than Amazon's sales.

1. Introduction

Over the last two and a half decades, Amazon has entered an increasing number of markets with its combination of product variety, low prices, and overall shopping convenience. Unlike Amazon, bricks-and-mortar stores — especially smaller ones — have limited capacity, are mostly limited to selling locally, and lack both data and data analytics. In this dire context, it is natural to ask whether there is any hope for the survival of traditional retail.

The purpose of our paper is to analyze the implications of Amazon's growth for the future of retail: Are brick and mortar stores doomed? If not, which ones are more likely to survive? And what strategic decisions can help them facing such a tough competitor? For instance, what type of products should they stock? These are some of the questions we address.

While these concerns — as well as our model — apply to virtually all retail industries, nowhere have they been more apparent than in the book retail market, Amazon's initial segment of choice. Accordingly, our analysis is motivated by and focused on the book-selling industry. That said, we believe our results have broader interest and applicability.

We consider a demand system with elements of horizontal differentiation (different book genres and different genre preferences) and vertical differentiation (different levels of book quality). Moreover, we assume that, all else equal, buyers have a preference for a specific channel (offline as opposed to online). Our model describes a bricks-and-mortar store's decision of whether to remain active and, if so, how to stock its shelves. We consider the trade-offs between a generalist bookstore and a specialist bookstore, i.e., one that is focused on a particular genre. Within the latter, we also distinguish between popular genres and niche genres. In various extensions of our baseline model, we consider the impact of pricing and exit decisions, competition between bricks-and-mortar stores, and consumer eclecticism.

Our central result is that, as Amazon becomes bigger (more available titles), a bookstore's optimal strategy is likely to shift from generalist to specialist. Intuitively, the store's choice trades off extensive margin, which favors a generalist approach, and intensive margin, which favors a specialist store. In other words, a generalist store attracts more potential customers, but a specialist store elicits greater willingness to pay from its patrons. As Amazon grows, the intensive margins of both generalist and specialist stores decrease equally. The generalist bookstore's extensive margin, by contrast, decreases at a faster pace than the specialist bookstore's extensive margin.

A series of additional results provide comparative statics with respect to key parameters. Specifically, for a given size of Amazon, smaller stores are more likely to follow a specialist strategy and more likely to survive. We thus predict a "polarization" of the firm-size distribution, with a large (online) player co-existing with multiple niche players and a declining number of mid-size and large bricks-and-mortar stores such as Barnes & Noble (Kahn and Wimer, 2019).

While this "vanishing middle" pattern has been observed by various authors in various contexts (e.g., Igami, 2011), our model also implies an additional, less obvious pattern: the bricks-and-mortar long tail. We show that, in equilibrium, bricks-and-mortar stores sell disproportionately more niche titles than Amazon. This runs counter to Chris Anderson's view of the Long Tail as it applies to online sellers (see also Brynjolfsson, Hu, and Simester, 2011):

> People are going deep into the catalog, down the long, long list of available titles,

> far past what's available at Blockbuster Video, Tower Records, and Barnes & Noble (Anderson, 2004).

Anderson's intuition is straightforward: Amazon's key advantage with respect to bricks-and-mortar stores is its lack of capacity constraints, which allows it to stock a large number of increasingly obscure titles. A bookstore that can only store one thousand books (for example) will instead focus, according to Anderson, on the most popular, mainstream titles. Why use precious and scarce shelf space on books that only attract few potential buyers?

What is missing from this observation and prediction is the endogenously-determined bricks-and-mortar store strategy, both in terms of size and especially in terms of specialization strategy. While it is true that an increasing fraction of total sales originates in niche products, our analysis suggests that this is *not* particularly true for online sellers; rather, a niche-oriented strategy and niche-genre sales are particularly expected from bricks-and-mortar sellers.

Interestingly, this implies that Amazon is responsible for two conceptually distinct long tails: its own, resulting directly from its virtually infinite catalogue; and an offline long tail, which is the byproduct of offline stores' specialization (which in turn is the response to Amazon's increasing dominance).

We then provide empirical evidence for our theoretical claims. By observing all sales made by a large publisher to different types of book retailers (independent bookstores, book chains, online retailers, airport bookstores) from 2016–2019 (a total of nearly six million transactions), we confirm that smaller independent bricks-and-mortar bookstores are more likely to follow a niche-genre strategy and, overall, show a thicker right-tail than Amazon's. By contrast, mass merchants such as Walmart do follow the pattern predicted by Anderson (2004), Brynjolfsson, Hu, and Simester (2011) and others.

■ **Road map.** The rest of the paper is structured as follows: we first review the existing literature. Section 2 contains our model, its main implications, and two main extensions (eclectic consumers and endogenous prices). Section 3 offers a discussion of the results, whereas in Section 4 we present our data and empirical findings (in the book market context). Section 5 concludes the paper.

■ **Related literature.** Conceptually, the paper that is closest to ours is probably Bar-Isaac, Caruana, and Cuñat (2012), who in turn build on Johnson and Myatt (2006). Bar-Isaac, Caruana, and Cuñat (2012) develop a model with a continuum of firms who set prices and choose their product design as general or specialized. Consumers, in turn, search for prices and product fit. Their main results pertain to the comparative statics of lower search costs, specifically how these lower search costs can lead both to superstar effects and long-tail effects. By contrast, our main focus is on the effect of an increase in a dominant firm's size (and quality, through better selection). Despite these differences, we share with Bar-Isaac, Caruana, and Cuñat (2012) the prediction that some firms "switch to niche designs with lower sales and higher markups" (p. 1142). An additional contribution with respect to Bar-Isaac, Caruana, and Cuñat (2012) is that, by considering the contrast between online and bricks-and-mortar stores, we illustrate the phenomenon of the bricks-and-mortar long tail, which departs from previous work, both theoretically and empirically.

Rhodes and Zhou (2019) observe that, in many retail industries, large sellers co-exist with small, specialized ones. They provide an explanation based on a model of consumer search frictions, showing that there exist equilibria where large, one-stop-shopping sellers co-exist with small, specialized sellers. We too provide an equilibrium explanation for the seller size distribution, albeit in a very different context (namely competition against a large online seller).

A number of authors have documented some of the patterns that motivate our analysis. Brynjolfsson, Hu, and Simester (2011) show that “the Internet channel exhibits a significantly less concentrated sales distribution when compared with the catalog channel.” This corresponds to the long-tail conventional wisdom as in Anderson (2004). In contrast, we argue theoretically and suggest empirically that the bricks-and-mortar long tail may actually be thicker than the online one.

Goldmanis et al. (2010) interpret the expansion of online commerce as a reduction in search costs and examine the impact this has on the structure of bricks-and-mortar retail. They examine data from travel agencies, bookstores and new car dealers and show that market shares are shifted from high-cost to low cost sellers. This is consistent with our theoretical predictions, though the mechanism is different.

Choi and Bell (2011) establish a link between the prevalence of preference minorities (consumers with unusual tastes) and the share of online sales. Using data from the LA metropolitan area, they find a strong link, even when controlling for multiple potential confounders. In similar vein, Forman, Ghose, and Goldfarb (2009) “examine the trade-off between the benefits of buying online and the benefits of buying in a local retail store,” and show that “when a store opens locally, people substitute away from online purchasing.” However, they “find no consistent evidence that the breadth of the product line at a local retail store affects purchases.”

Consistent with both our theory and recent anecdotes from the US book market, Igami (2011) conducts an empirical analysis of Tokyo’s grocery market and finds that the rise of large supermarkets does not crowd out small, independent stores, but rather mid-size ones. Furthermore, we suggest that niche specialization — a strategy not available to (or at least not optimal for) mid-size retailers — is an important driver of small stores survival, suggesting that these results might fail to hold in markets in which specialization is not a possibility in the first place.

Neiman and Vavra (2019) observe that “the typical household has increasingly concentrated its spending on a few preferred products.” They argue that this is not driven by “superstar” products, rather by increasing product variety. “When more products are available, households select products better matched to their tastes.” They also argue that the distinction between online and offline sales does not play an important role in explaining this trend.

Focusing on the US book market, Raffaelli (2020) summarizes the drivers of independent bookstores’ recent success in three Cs: curation (“Independent booksellers began to focus on curating inventory that allowed them to provide a more personal and specialized customer experience”), convening (“Intellectual centers for convening customers with likeminded interestz”) and community. All of these strongly resonate with both our theoretical and empirical findings.

Table 1
Main notation used in the paper

Variable	Description
a, b	online store (Amazon) and offline (bricks-and-mortar) stores
k, c	store b's capacity and cost per unit of capacity
$\tilde{d}, d$	horizontal distance from bricks-and-mortar store b_0; $d = \max \tilde{d}$
f, F	pdf and cdf of $\tilde{v}$
g, s	general and specialty store
$m(t)$	maximum $\tilde{v}$ from t draws out of $F(\tilde{v})$
p	book price
q	bricks-and-mortar's store market share
t	number of titles
$\tilde{v}, \tilde{w}$	vertical and horizontal preferences (maximum values: v and w)
x, y	popular and niche genre
z	total number of titles (carried by store a)
α, β	popularity of x, fraction of b's capacity devoted to$x$
π	store b's profit
τ	transportation cost (when b_0 and b_1 compete)

2. Theory

Consider an economy with two book sellers, a (Amazon) and b (bricks-and-mortar); and two different book genres, x and y. (Considering the large number of different variables used in the paper, Table 1 lists the main notation used in the paper.) While we assume there exists only one bricks-and-mortar store, our intent is to model this as a generic bricks-and-mortar store, assuming that its effective competitor is the online store. Later we also consider the possibility of competition between bricks-and-mortar stores.

We assume that there is a measure one of book buyers, equally split into two types, x lovers and y lovers.[1] Buyers of type x (resp. y) have a value v for one book of genre x (resp. y) and zero for any book of genre y (resp. x), where the value of $\tilde{v}$ is generated from a cdf $F(\tilde{v})$, where $f(\tilde{v}) > 0$ if and only if $\tilde{v} \in [0, v]$, where v is possibly infinite.[2]

We assume that, independently of preferences for x and y, book buyers have a preference for firm b (with respect to firm a). This may reflect an intrinsic taste for in-person shopping, the presence of additional amenities (which we endogenize in one of our extensions), a desire to support small and local businesses, or an ideological aversion to (or taste for) Amazon.[3]

1. Later in the paper, we consider the asymmetric case, that is, the case of a popular genre and a niche genre.
2. Later in the paper, we extend this to the case of eclectic consumers, who have positive valuations for both genres.
3. Saxena (2022) describes recent examples of independent bookstores providing offline perks such as bars and cafes.

We assume that this preference (extra welfare from buying at a local store), denoted by $\tilde{w}$, is uniformly distributed in $[0, w]$.[4]

Seller a carries all titles in the economy, a total of z titles, $z/2$ of each genre. By contrast, seller b can only carry k titles, that is, k measures the bookstore's capacity. Book prices are constant and exogenously given (until later in this section), and with no further loss of generality we assume prices are equal to \$1.

At a given seller, buyers can learn both the genre and the value $\tilde{w}$ of a title at no cost. By contrast, when b chooses what books to carry, it can observe genre but not $\tilde{v}$. Therefore, the bookstore determines which type of books to sell but otherwise selects a random sample of values $\tilde{v}$. Each buyer selects the bookseller providing the highest expected value and, within a given bookstore, buys the one book that yields the highest value $\tilde{v}$. If the store carries t titles of the buyer's preferred genre, then the buyer receives an expected value $m(t)$, where $m(t)$ is the expected value of the highest element of a sample of size x drawn from $F(\tilde{v})$.

■ **General or specialty store?** The focus of our analysis is on bookstore b's strategy as the value of z increases. Specifically, firm b (the bricks-and-mortar store) has three options: to exit, to remain active as a general store, and to remain active as a specialty store. A general store carries $k/2$ titles of each genre, whereas a specialty store carries k titles of a given genre.

We first consider the case when b pays no fixed cost to remain active, so that it's a dominant strategy to do so. The only question is then how to design the store, namely whether to be a general or a specialty store. We present our results both as comparative statics with respect to the value of z (a measure of the online store's growth), and k (the bricks-and-mortar store's capacity). Our first two results are based on the following assumptions:

Proposition 1. *Suppose that*

$$w < \min \{m(z/2), v - m(k/2)\}$$

(a) There exists a threshold $z_{gs} = z_{gs}(k, w)$ such that an active firm b optimally chooses to be a specialty store if and only if $z > z_{gs}$. Moreover, $z_{gs}(k, w)$ is increasing in both k and w.
(b) There exists a threshold $k_{gs} = k_{gs}(z, w)$ such that an active firm b optimally chooses to be a specialty store if and only if $k < k_{gs}$. Moreover, $k_{gs}(z, w)$ is increasing in z and decreasing in w.

The proof for this and all other results can be found in the Appendix. In order to understand the intuition for Proposition 1, note that the choice between a general and a specialty store trades off an "extensive margin" and an "intensive margin" effect. By switching to a specialty strategy, a store forgoes half of its potential customers, those interested in the genre that is no longer stocked (extensive margin). On the other hand, by stocking twice as many titles of a given genre, the store increases the expected quality that a patron expects from visiting

4. The assumption that the lower bound of $\tilde{w}$ is zero simplifies the analysis and is without loss of generality. That is, all of our results would be unaffected if we assumed a negative lower bound for $\tilde{w}$, corresponding to a relative preference for firm a. The reason for this is that, because Amazon has a size advantage ($z > k$), a positive $\tilde{w}$ is required to buy offline. Put differently, all consumers with $\tilde{w} < 0$ or $\tilde{w} = 0$ purchase from Amazon, so that we can simply assume $\tilde{w} \geq 0$ when determining the critical z threshold.

the store (intensive margin). As total supply z increases, the expected payoff from visiting store a, $m(z)$, increases.

As z increases, store a becomes relatively more attractive, which in turn lowers the demand for store b. This increase in valuation for store a hurts the general store b more than the specialty store b. Basically, the general store loses readers from both genres, whereas the specialty store only looses readers from a smaller set. It follows that, starting from a point where a general store strategy is better, there exists a threshold value of z past which a specialty store strategy yields higher profit.

Another way of understanding Proposition 1 is that, as z increases, the profit of both a general and a specialty store decrease. However, the profit of a general store decreases at a faster rate. In other words, specialty stores are better "insured" against Amazon's growth, whereas general stores — such as Barnes & Noble or the now defunct Borders — are likely to suffer bigger profit losses.

The condition in Proposition 1 ensures that the solution is interior. If the condition does not hold, then we are in a corner solution whereby it is a dominant strategy for b to be a general store.

We consider comparative statics in both k and w. First, for a given value of z, a store with larger capacity k is less likely to specialize, that is, it requires a larger Amazon for such a store to abandon a generalist strategy. Or, to put it differently, store b's decision to specialize is based on its *relative* size with respect to Amazon.[5] Similarly, the threat posed by Amazon is lower the greater w, that is, the greater the buyers' aversion to purchasing from Amazon. Accordingly, given z and k, store b is less likely to become a specialty store as a strategy to cope with online competition the higher w is.

Industry players understand these dynamics. James Daunt, CEO of UK chain Waterstones, argues that

> [Amazon's] unmatchable scale is liberating for booksellers; it means stores can focus on curating books that communicate a particular aesthetic, rather than stocking up on things people need but don't get excited about (Todd, 2019).

In private communication, Mark Cohen, Director of Retail Studies at Columbia GSB, echoes this view:

> There is a tremendous resurgence of local bookstores, but these have relevance because (...) they're not trying to be all things to all people as Barnes & Noble has always tried to be. They're either picking on a genre or curating an assortment that appeals to a local customer.

Part (a) of Proposition 1 highlights the dynamic interpretation of online-offline competition, namely what happens as Amazon increases in size; conversely, part (b) highlights the cross-sectional interpretation, namely what happens to large and to small bricks-and-mortar stores. The data we will analyze in Section 4 comprises thousands of stores but only a handful of years. For this reason, while parts (a) and (b) of Proposition 1 are essentially equivalent, we will focus on the cross-sectional implications.

■ **Niche genres.** So far we have assumed that both genre x and genre y have the same popular appeal. A more realistic case has one of the genres — say, genre x — be a popular

5. Non-linearities in $m(\cdot)$ imply that the ratio k/s is not a sufficient statistic for the specialization decision. Nevertheless, the specialist strategy is more likely when either k is small or z is large.

Figure 1
Bookstore profits from specializing in popular genre (π_x) or niche genre (π_y) as a function of z when $F(\tilde{v}) = \tilde{v}/v$.

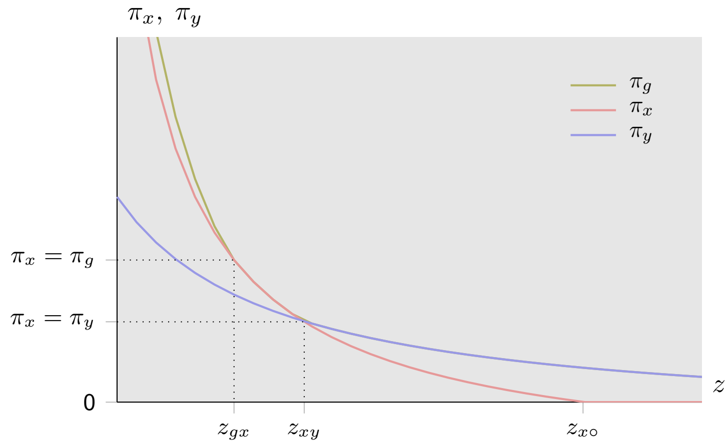

genre, whereas y is a less popular one — a niche genre. Suppose that there is a measure 1 of potential book buyers, α of which are only interested in genre x books; and suppose that $\alpha > \frac{1}{2}$. (So far, we have implicitly assumed that $\alpha = \frac{1}{2}$.) Consistent with the assumption that genres x and y have different popular appeal, we assume that a fraction $\alpha\, z$ of the total titles are of genre x, and a fraction $(1-\alpha)\, z$ are of genre y.

Proposition 1 states that, as z increases, store b optimally switches from general to specialty store. The next proposition complements that result by stating that, within the specialty strategy, store b optimally chooses the niche strategy if z is high enough.

Proposition 2. *Suppose that*

$$w < \min\left\{m(z/2), v - m(k/2)\right\}$$

(a) There exists an z_{xy} such that an active store b specializes in a niche genre (rather than a popular genre) if $z > z_{xy}$.
(b) There exists a k_{xy} such that an active store b specializes in a popular genre if $k > k_{xy}$.

Figure 1 illustrates Proposition 2. The key insight is that, *relatively* speaking, a niche-genre store suffers less from an increase in z than a popular-genre store, in a way that is similar to, but different from, the general-specialist trade-off considered in Proposition 1. For low values of z, the advantage of a niche-genre store, in terms of higher intensive margin, is outweighed by the simple fact that a popular genre is more popular, that is, attracts a greater number of potential customers. For high values of z, however, the niche strategy becomes increasingly attractive, as illustrated by Figure 1. Specifically, for $z > z_{xy}$, π_y, the profit from a niche-genre strategy, is greater than π_x, the profit from a popular-genre strategy.

Formally, the proof of Proposition 2 proceeds by deriving the value z_x when $\pi_x = 0$ and establishing that, at that value, $\pi_y > 0$. This proof strategy is similar to that of Proposition 1. There is one difference, however. In Proposition 1, we show that $z > z_{gs}$ is a necessary and sufficient condition for specialization. By contrast, in Proposition 2 $z > z_{xy}$ is only a

Figure 2
Optimal stocking policy for generalist store (assuming v is uniformly distributed). α is the fraction of genre x buyers, whereas β is the fraction of genre x books optimally stocked by a generalist store.

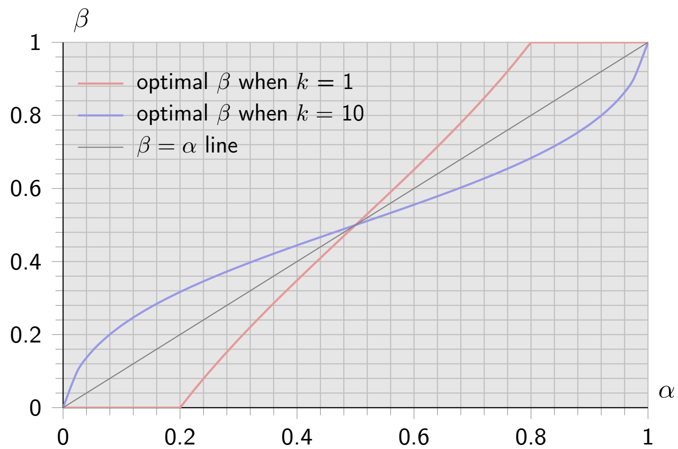

sufficient condition. The difference stems from the fact that we can prove the monotonicity of $\pi_s - \pi_g$ in general terms but not the monotonicity of $\pi_y - \pi_x$. If we further assume that v is uniformly distributed, then the condition $z > z_{xy}$ becomes a necessary and sufficient condition.[6]

An implication of this result is that bricks-and-mortar sales are more niche-concentrated than online sales (or total sales). In other words, we uncover a novel reason why Amazon is leading (indirectly) to a thickening of the long tail. We return to this in the next section.

■ **General, popular-genre, and niche-genre stores.** A natural extension of the analysis so far is to integrate the choice of generalist vs specialist (Proposition 1) with the analysis of genre of specialization (Proposition 2). In our initial model, we assumed two equal-sized genres, x and y. In this context, a general bookstore is one that stocks x and y in equal amounts, whereas a specialty bookstore is one that stocks either only x or only y. When there are two genres of different sizes, as in the model underlying Proposition 2, the decision of how to stock is not trivial. Suppose that a fraction α of the titles (and a fraction α of the potential demand) correspond to genre x. Let β be the fraction of a general store that carries genre x books. Should β be greater than, equal to, or lower than α?

Figure 2 illustrates this decision in the case when $F = v$, and so $m(t) = t/(1+t)$. If the value of k is small ($k = 1$ in the present example), then the optimal stocking policy is to over-stock the most popular genre. This is shown by $\beta > \alpha$ for $\alpha > \frac{1}{2}$ (red line). By contrast, if the value of k is large ($k = 10$ in this example), then the optimal stocking policy is to over-stock the least popular genre. This is shown by $\beta > \alpha$ for $\alpha < \frac{1}{2}$ (blue line). Intuitively, when k is large, then the marginal value of an extra title is lower, due to concavity of $m(k)$. This is particularly true for a popular genre. Therefore, in relative terms and at the margin, the seller is better off by stocking a title of a niche genre. By contrast, if k is small, then the extensive margin effect dominates and the seller is better off

6. The proof can be obtained from the authors upon request.

by overstocking (relatively speaking) the popular genre.

Taking into account the optimal stocking strategy, Figure 1 plots the profit of a general store (as well as the profit function of a specialty store focused on a popular genre, x, or on a niche genre, xy). As can be seen, as z increases, firm b's optimal choice shifts from being a general store to being a specialty store focused on the popular genre to finally being a specialty store focused on the niche genre. In this way, Figure 1 illustrates both Proposition 1 and Proposition 2.

■ **Exit.** Suppose now that the bricks-and-mortar store must pay a fixed cost $c\,k$ in order to operate, where c is cost per unit of capacity. That is, let profits be given by

$$\begin{aligned} \pi_g(z,k) &= \left(1 - \left(\frac{m(z/2) - m(k/2)}{w}\right)\right) - c\,k \\ \pi_s(z,k) &= \tfrac{1}{2}\left(1 - \left(\frac{m(z/2) - m(k)}{w}\right)\right) - c\,k \end{aligned} \tag{1}$$

Now that we assume $c > 0$, a third option — exit — becomes non-trivial. We consider the bookstore's optimal choice in the (z,k) space, now a choice between being a general store, a specialty store, or simply exiting. (We return to assuming two genres of equal size, so that the only relevant decision is whether to specialize, not what genre to specialize in.)

From Proposition 1, we know that there exists a threshold z_{gs} such that, conditional on being active, a specialty-store strategy is better than a general-store strategy if and only if $z > z_{gs}$. Equivalently, there exists a threshold k_{gs} such that, conditional on being active, a specialty-store strategy is better than a general-store strategy if and only if $k < k_{gs}$. We are now interested in characterizing the joint exit and specialization behavior by bookstores of different sizes as Amazon size grows increasingly large.

Proposition 3. *Suppose that*

$$\pi_g\big(z'', k_{gs}(z'', w)\big) > 0 \tag{2}$$

where $\pi_g(\cdot,\cdot)$ is given by (1) and $k_{gs}(\cdot,\cdot)$ is given by Proposition 1. Then there exist $0 < k_1 < k_2 < k_3 < k_4 < \infty$ such that, as z increases from z' to $z'' > z'$,

- *If $0 < k < k_1$, then a specialist store remains a specialist*
- *If $k_1 < k < k_2$, then a generalist store becomes a specialist*
- *If $k_2 < k < k_3$, then a generalist store remains a generalist*
- *If $k_3 < k < k_4$, then a generalist store exits.*

Proposition 3 may be summarized by stating that, as z increases, bricks-and-mortar stores become predominantly specialist stores. This results from two forces: First, large, generalist stores are more likely to exit than small, specialist stores. Second, some generalist stores convert to specialist stores. In other words, Proposition 3 provides a formal description of the "vanishing middle" phenomenon. In terms of our model, this mens that there is an increase in sales by the very large seller a and by an increasing number of small, specialized b bricks-and-mortar sellers.

Proposition 3 requires (2) to hold. This condition implies that we are not in a corner solution such that all generalist stores prefer to exit. In other words, we provide conditions such that the "middle" (large bricks-and-mortar stores) vanishes gradually, not suddenly.

Figure 3
Optimal choice in the (z, c) space (number of titles, fixed cost) in the linear case

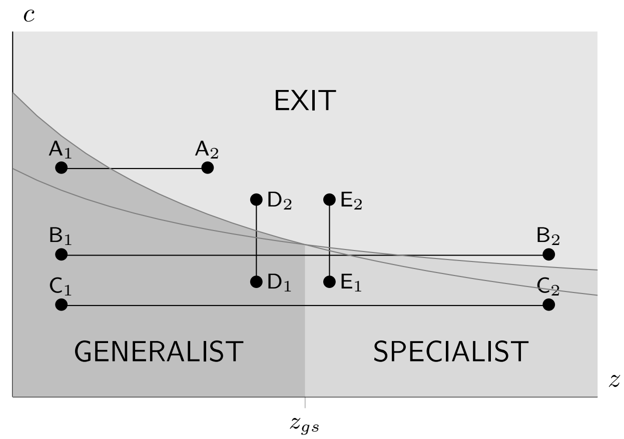

Our result is in line with Igami's (2011) empirical findings. He studies the competitive effects of large, one-size-fits-all supermarket entry in Tokyo in the early 2000s. He finds that it is not the smallest, specialized food sellers that exit, but rather the mid-sized, more generalist ones.

Proposition 3 characterizes the bricks-and-mortar store choice in the (z, k) space, where z is Amazon's size and k the store's size. Alternatively, we can also look at industry evolution in the (z, c) space, where c is the cost of one unit of capacity. Figure 3 does so in the linear case, that is, $F(\tilde{v}) = \tilde{v}/v$. The boundary $c^\circ(s)$ is the minimum of two boundaries, the exit boundary for a general store and the exit boundary for a specialty store, both of which are plotted in Figure 3. Together with the z_{gs} threshold, these lines define three regions: the GENERALIST region, defined by $z < z_{gs}$ and $c < c^\circ$ (effectively, the generalist exit boundary); the SPECIALIST region, defined by $z > z_{gs}$ and $c < c^\circ$ (effectively, the specialty exit boundary); and the EXIT region, defined by $c > c^\circ$.

It's unlikely that there have been any major changes in the fixed cost of keeping a bricks-and-mortar store open (except for the general increase in commercial real estate prices in some areas). Aside from Amazon, the most relevant changes in terms of the cost and benefit of operating a store in a given location are likely to be related to local demographics. In our model set up, we normalize price and quantity per title. As such, the relevant changes in demographics are absorbed in the value of the fixed cost $c\,k$. So, for example, an increase in income in a given neighborhood would be measured by our model as a decrease in c. In what follows, we consider this interpretation of the value of c.

Based on Figure 3, we may consider several possible exogenous changes in z and c. Moves A, B and C correspond to an increase in the number of titles, z. In case A, we have a store with a high value of c, which we may interpret as a neighborhood with demographics unfavorable to book selling. As the value of z increases, we observe a general store exit. (Recall that, if z is small enough, then all stores are general stores.) In other words, considering the store's relatively low "efficiency" (as measured by c) the store does not even try the strategy of being a specialty store, it simply cannot put up with a's competition.

By contrast, in case B we have a store with a lower value of c. As with store A, B starts off as a general store when z is low. As z increases, long after store A has gone out of business, B remains active, but past $z = z_{gs}$ becomes a specialty store. As z continues to increase, B eventually exists as well.

Finally, in case C we observe a store that is sufficiently efficient (in the sense of having a low value of c) that, no matter how high z is, it remains active. Notice however that, similarly to B, store C becomes a specialty store when $z > z_{gs}$.

Moves D and E correspond to a increase in c. In case D, we observe the exit of a general store, whereas in case E we observe the exit of a specialty store. As mentioned earlier, a change in c is best interpreted as a change in local demand conditions (since c is effectively measured in units of consumer demand).

Figure 3 also helps understand the contrast between urban and suburban/rural areas. If a bricks-and-mortar store has limited spatial reach, then it makes sense to think of urban areas as areas where each store has a higher potential demand, which in turn corresponds to a lower value of c. One might argue that urban density also implies higher costs, in particular real-estate costs. However, if the long-run supply of real estate is relatively flat, then an increase in density leads to an increase in the ratio of density over monetary cost, which effectively corresponds to a lower c.

Now suppose that the value of of z is close to the disruption level z_{gs}. Suppose moreover that, empirically, store heterogeneity within a certain area corresponds to variation in c and, in particular, variation in the effective value of z for that store. For example, there might be variation in store-specific preference which enters the profit function in the same way as a variation in z does. In this context, as we compare an urban area (low value of c, something like level C in Figure 3) with a suburban area (high value of c, something between levels A and B in Figure 3), we observe that, in the former, stores are either general of specialty stores; whereas, in the latter, they are either general stores or exiters. This implies that, starting from a a certain distribution of general and specialty stores, we would expect the distribution of stores in the urban area to skew in the direction of specialty stores. We provide empirical support for this conjecture in Section 4.

It is important to note that this relation between market density and the skew toward specialization is *not* due to the classical Adam Smith argument that the division of labor is limited by market size. In fact, moving along a vertical line (cases D and E in Figure 3) does not change the degree of specialization, only the entry/exit decision. Our point is that the combination of entry/exit decisions and the disruption caused by changes in z may lead to an observed association between market density and specialization even if we assume constant returns to scale.

■ **Endogenous prices.** So far, we have assumed that all books are priced \$1. This has allowed us to focus on the main issues regarding specialization while keeping the analysis tractable. We now explicitly consider pricing choices. Our goal is to verify the robustness of our previous findings as well as to develop additional intuition regarding the comparative statics of Amazon's expansion.

Recall that the actual market structure we have in mind includes one dominant firm and a large number of fringe firms. Although for simplicity we focus on the decisions of one representative fringe firm, it makes sense to treat firms a and b as different types of strategic players. Consistent with with this interpretation, we assume that firm a acts a price leader by setting p_a first.

Given p_a, the bricks-and-mortar store b responds by setting its price, which we denote by

p_g if the store is a general store and p_s if the store is a specialty store. Our focus in on firm b's decisions. Accordingly, we take p_a as an exogenous variable (and later consider comparative statics with respect to it).[7] Similar to Propositions 1 and 2, we make a parameter assumption so as to eliminate trivial corner solutions (if the following assumption fails to hold, then we may be in the case in which a specialty store is always optimal).

In what follows, we first solve for store b's optimal price and then reconsider the store's optimal positioning (general or specialty). Our next result extends the main intuition of Proposition 1, adding one new dimension of comparative statics.

Proposition 4. *Suppose that*

$$p_a > w + \frac{m(k) - \sqrt{2}\, m(k/2)}{\sqrt{2} - 1}$$

There exists a threshold z_{gs} such that store b optimally chooses to be a specialty store if $z > z_{gs}$. In the right neighborhood of z_{gs}, the specialty store sets a higher price, captures a lower market share and earns a higher profit than a general store.

When discussing Proposition 1, we argued that the trade-off between a general and a specialty store is a trade-off between the extensive margin (which favors a general store) and the intensive margin (which favors a specialty store). The proof of Proposition 4 establishes that, when it comes to price setting, only the intensive margin matters. This explains why a specialty store sets a higher price than a general store. By devoting its space to one book genre only, a specialty store elicits a higher willingness to pay from buyers interested in that genre, which in turn allows the store to set higher prices. This in turn increases the store's incentives to specialize.

Similar to Proposition 1, Proposition 4 establishes that, if firm a is big enough (high z), then firm b is better off by becoming a specialty store. The main intuition for the z-threshold part of Proposition 4 is similar to Proposition 1: As total supply z increases, the specialty store option becomes *relatively* more attractive. In sum, the first part of Proposition 4 shows that the intuition from Proposition 1 is robust to the introduction of pricing.

The novel aspect of Proposition 4 is its second part, the statement that, past the disruption level z_{gs}, a specialty store sets a higher price, captures a *lower* market share and earns a higher profit than a general store. We call this the *boutique effect*. The specialty store in the model with fixed prices trades-off extensive margin and intensive margin so as to maximize the number of customers. By switching from general to specialty store, firm b loses potential customers, but its offering becomes so much more attractive to its reduced set of customers that it ends up attracting more customers. By contrast, once we introduce prices we observe that the switch to a specialty-store strategy not only sacrifices potential demand but also sacrifices actual demand. Such drop in actual demand is more than compensated by an increase in the intensive margin via higher sale prices.

■ **Eclectic consumers.** So far we have assume that consumers are divided into x fans and y fans. Specifically, the value v of a book outside of a consumer's preferred genre is zero. At the opposite extreme, consider the case when consumers are totally eclectic, that is, they value both genres equally.

7. Endogeneizing Amazon's decisions would be a promising direction for future research.

Clearly, eclectic consumers are bad news for specialty stores. Before, an x fan valued a specialty store at $m(k)$ and the online store at $m(z/2)$. By contrast, an eclectic consumer values the online store at $m(z)$ whereas the specialty store is still valued at $m(k)$ (here we are excluding the preference parameter z).

Regarding a general store, the analysis is not as obvious. Before, the value of a general store was $m(k/2)$ for an x fan or a y fan, whereas the value of the online store was $m(z/2)$. By contrast, an eclectic consumer values the online store at $m(z)$ whereas the general store is valued at $m(k)$ (again, we are excluding the preference parameter z). In which case is the general store better off? The answer depends on which difference is greater, $m(z/2)-m(k/2)$ or $m(z)-m(k)$. Notice that $m(z)-m(k) > m(z/2)-m(k/2)$ if and only if $m(z)-m(z/2) > m(k) - m(k/2)$. Since $z > k$, $z - s/2 > k - k/2$, which would suggest the inequality holds. However, concavity of $m(t)$ would work against the inequality. Suppose that $F = v$ is linear, so that $m(t) = t/(1+t)$. Then the function $m(t)-m(x/2)$ is non-monotonic, first increasing for $x \in [0, \sqrt{2}]$ and then decreasing. This implies that we can find values of z and k such that the inequality is in turn true or false. So, even assuming a specific distribution of v, we cannot guarantee that a general store is better off or worse off when serving eclectic consumers rather than polarized consumers.

It has long been argued that Amazon benefits from increased consumer specialization, and that this is largely the purpose of its recommendation system: by presenting each consumer with increasingly personalized offerings, it makes bookstores obsolete, since bookstores cannot, due their limited size, cater to each consumer's idiosyncrasies. However, as the above analysis shows, this is not necessarily true when we endogenize bricks-and-mortar stores' strategies: more specialized consumers allow specialty stores to emerge, which can be detrimental to Amazon's profits.

■ **Offline Amenities.** What can bookstores do when consumers are eclectic? We know that, *ceteris paribus*, bookstores' survival is crucially dependent on the relative consumer preferences for offline shopping.

While so far we have treated this distribution as exogenous, it is interesting to consider the case in which bookstores explicitly invest in it, for instance by boosting their distinctly offline, or "social", features: readings, cafes, bars, but also personalized staff recommendations, for instance. These features are appealing in that they can not be directly replicated by Amazon. They are also increasingly widesparead: see Raffaelli (2020) for a discussion of "community" as one of the pillars of brick-and-mortar bookstores survival, and Saxena (2022) for some recent examples.

Are offline amenities a complement or a substitute of specialization? And which stores benefit the most from them? We have the following:

Proposition 5. *$\partial^2 \pi / \partial w \, \partial z > 0$ and $\partial^2 \pi / \partial w \, \partial k < 0$, where π stands for either π_g or π_s. Moreover, $\partial \pi_g / \partial w > \partial \pi_s / \partial w > 0$ and $d^2 \pi_g / dw^2 < d^2 \pi_s / dw^2 < 0$.*

In words, $\partial^2 \pi / \partial w \, \partial z > 0$ states that the benefits from improved amenities, which we model by an increase in w, is increasing in the value of z. So, the greater the size of the online store, the greater the incentive for bricks-and-mortar stores to invest in amenities that shift consumer preference in favor of bricks-and-mortar stores. The "time series" comparative static (increase in z) is complemented by a "cross section" comparative static (variation in

k). In this case, Proposition 5 states that smaller stores have a greater incentive in investing in amenities than larger stores.

The second part of Proposition 5 compares incentives in terms of store type rather than store size. It states that a generalist store has a greater incentive to invest in increasing w, but the returns from such an investment are themselves decreasing. The idea is that investments in w appeal to all consumers. Therefore, they are more valuable to those stores (generalist stores) which did not give up on half of the consumers to begin with. Moreover, amenities compensate for a "quality gap" in catalogue terms, and this gap is larger for generalist stores than for specialty stores, that is, $m(z/2) - m(k/2) > m(z/2) - m(k)$. Finally, the fact that amenities offer decreasing returns follows straightforwardly from the concavity of the demand function with respect to w.

Proposition 5 suggests that, for a small store (low k), investing in amenities (i.e., increasing the value of w) may provide an alternative strategy to specialization. This is particularly the case when a significant fraction of consumers are eclectic (so that the gain from becoming a specialist is not as big).

Finally, we note that, by Proposition 4, higher prices correspond to a higher w. Thus, similar to specialization, improving offline amenities allows bricks-and-mortar stores to charge higher markups. Unlike specialization, however, this conclusion is robust to different preference specifications for consumers.

■ **Bricks-and-mortar store competition.** Up to now, we considered competition between one online store and one bricks-and-mortar store. Implicitly, the idea is that there are a plethora of small (possibly independent) bricks-and-mortar stores with a catchment area that does not overlap with any other bricks-and-mortar store. Consider now the case when two bricks-and-mortar stores, say b_0 and b_1, do compete for the same potential demand. Specifically, we assume a consumer is characterized by a bricks-and-mortar-store preference $\tilde{w}$ and a relative preference between stores b_0 and b_1 in the form of a location $\tilde{d} \in [0, 1]$ and transportation cost τ per unit of distance to store b_0 (located at 0) and to store b_1 (located at 1). Moreover, we assume that $\tilde{d}$ and $\tilde{w}$ are independently and uniformly distributed: $\tilde{d} \sim U[0, 1]$ and $\tilde{w} \sim U[0, w]$. Our main result is that, under competition, the choices of genre by stores b_0 and b_1 exhibit strategic complementarities.

Proposition 6. *Let z be such that store b_0 and b_1 are indifferent between a general- and a specialty-store strategy absent offline competition. In the neighborhood of z, being a specialty store is a strict best response to the rival choosing to be a specialty store.*

Proposition 6 suggests that competition provides an additional force pushing in the direction of specialization. Suppose that we fix firm b_1's strategy at being a general store. As z crosses a certain threshold, say $z_\circ$, then firm b_0's optimal strategy switches to becoming a specialty firm (of either x or y). However, if firm b_1 has become a specialty firm (choosing, say, genre y), then, *even if z is lower than $z_\circ$* (by a little), firm b_1 also optimally switches to being a specialist (specializing in the niche that firm b_1 did not).

To conclude this section, we note how Amazon is strictly worse off when competing with two specialty stores compared to two generalist stores, as the overlap between the latter is far greater than between the former. Again, this suggests caution when interpreting a higher degree of consumer specialization as a desirable outcome for larger, online retailers.

Figure 4
Firm profit and consumer welfare. Effects of switching from general to specialty x store.

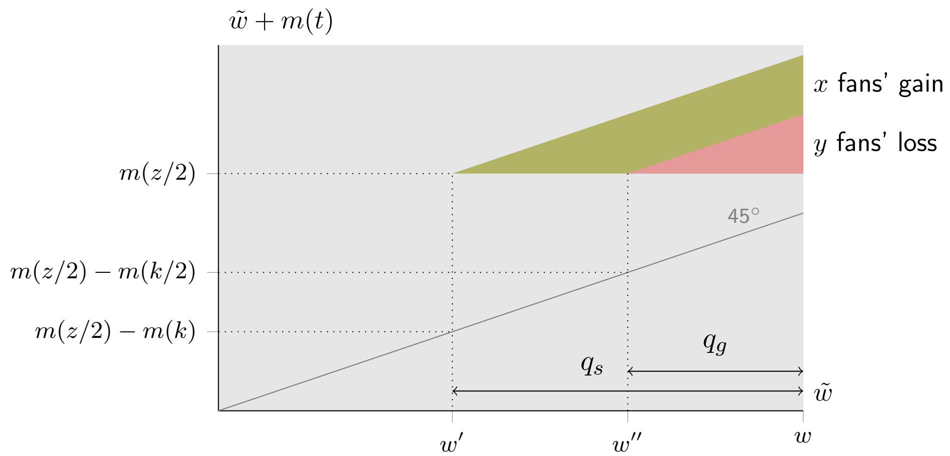

■ **Welfare analysis.** All of our analysis so far has focused on firm b's profits and optimal choices. A natural follow-up question is the relation between firm b's decisions and consumer welfare. Let us go back to the model with fixed prices and one bricks-and-mortar store, firm b. Let us consider, as in the initial model, the choice between being a general and being a specialty store. Suppose social welfare is given by consumer surplus plus firm profits. Since all sellers set $p = 1$ and the market is covered (all consumers make a purchase), consumer surplus is a sufficient statistic of social welfare.

Figure 4 illustrates the contrast between a general and a specialty store when competing against firm a. On the horizontal axis we measure each consumer's value of z, that is, their disutility from buying from firm a. On the vertical axis we measure the advantage, in terms of vertical quality, of the online store with respect to the bricks-and-mortar store. The 45° line measures the points at which the "horizontal" differentiation advantage of firm b exactly compensates the "vertical" differentiation advantage of firm a.

Consider first the case of a general store b. Its disadvantage with respect to store a is given by $m(z/2) - m(k/2)$. It follows that only consumers with a value of $\tilde{w}$ greater than w'' purchase at the bricks-and-mortar store. Since $\tilde{w}$ is uniformly distributed, we conclude that firm b's market share is given by $q_g = w - w''$.

Consider now the case of a specialty store b. Its disadvantage with respect to store a is given by $m(z/2) - m(k)$. It follows that only consumers with a value of $\tilde{w}$ greater than w' purchase at the bricks-and-mortar store. Since $\tilde{w}$ is uniformly distributed, we conclude that firm b's market share (among its genre followers) is given by $q_s = w - w'$. However, we must keep in mind that if firm b focuses on genre x, for example, then it loses potential buyers who are only interested in y. In other words, by becoming a specialty store firm b halves its potential demand. Therefore, its market share is $(w - w')/2$.

The values of z and k were selected so that $\pi_g = w - w'' = (w - w')/2 = \pi_s$. In other words, for the particular values of z and k underlying Figure 4, firm b is indifferent between being a general store or being a specialty store. Consumers, however, are not indifferent between the two types of store. Consumer surplus is given by the area below

$$\max\{m(z/2), \tilde{w} + m(\widetilde{k})\}$$

where $\widetilde{k} = k/2$ or $\widetilde{k} = k$ for a general and a specialty store, respectively. It follows that, for genre x consumers, the switch from a general to a genre x specialty store implies an increase in consumer surplus given by the green trapezoid in Figure 4. By contrast, for genre y consumers the switch implies a decrease in consumer surplus given by the red area in Figure 4. By construction, the green area is greater than the red area. More generally, we have just established the following result:

Proposition 7. *When store b is indifferent between being a general or a specialty store, the average consumer strictly prefers the latter.*

Intuitively, consumer surplus is "convex" in the vertical utility provided by the bricks-and-mortar store. This implies that consumers prefer the "bet" of having a specialty store of their preferred genre with probability 50% than a general store with probability 100%.

This intuition is related to a number of results in the IO literature. Mankiw and Whinston (1986) provide conditions such that, in equilibrium, there is excess entry into a market. Intuitively, the entrant does not correctly take into account the positive externality it creates for consumers nor the negative externality it creates for its competitors. Similarly, our firm b does not take into account the positive surplus effect it has on the consumers who like the genre in which they specialize.

3. Discussion

We believe that our theoretical findings have important practical implications for marketing and strategy, namely in the context of bookstores and other retail markets. In this section, we discuss some of these implications.

■ **Barnes & Noble.** In 2019, Barnes & Noble appointed James Daunt as its new CEO. Daunt was previously the founder of Daunt Books and managing director of UK's large bookshop chain Waterstones (Chaudhuri, 2019). Daunt's philosophy, as he puts it, is centered around some core tenets (Segal, 2019):

- Escape broad genres, such as "self-help" or "history", organizing bookstores around some specific, and often niche, themes;
- Curate selections locally, allowing the local staff to pick books, and avoiding general, UK-wide catalogs;
- Avoid the convenience trap, focusing on the many perks of the offline experience instead.

This business strategy resonates with our theoretical findings. First, and most obvious, Daunt clearly emphasizes the importance of specialization (Proposition 1 and 2), thus avoiding broad genres on which Amazon's advantage is hard to counteract.

That said, it is important to note that for this form of bricks-and-mortar specialization to arise, a substantial fraction of consumers need to be specialists, that is, have genre-specific preferences. When consumers are eclectic, offering offline amenities such as readings and cafes, but also curated staff recommendations, may prove a more fruitful strategy (Proposition 5), as also highlighted by Daunt.

■ **The tyranny of the majority.** In his influential book, Waldfogel (2007) states that

> When fixed costs are substantial, markets provide only products desired by large concentrations of people.

Our analysis suggests that the competition between an ever-larger online platform and bricks-and-mortar stores may actually counter Waldfogel's "tyranny of the majority." In other words, while we acknowledge that there is empirical evidence for Waldfogel's prediction, we argue that Amazon's increased dominance might have at least partly reversed this picture in a variety of retail markets. Chief among them is arguably the book market, which combines early Amazon penetration with enormous product variety.

Nevertheless, our results have some parallels to Waldfogel's. Proposition 7 shows that the extent of specialization is insufficient: consumer welfare would increase with more specialization than it results in equilibrium.

■ **Amazon's embarrassment of niches.** Amazon's highly personalized algorithms have long been believed to fracture consumers into taste niches, lengthening the tail in sales and thus the value of Amazon's virtually infinite inventory. Our analysis highlights a potential drawback to Amazon's strategy: as more consumers acquire (or discover) a specific taste, smaller retailers respond by targeting these increasingly relevant taste communities. In other words, taking into account bricks-and-mortar specialization decisions, it is unclear whether consumer specialization is good news for Amazon after all.

■ **A contrast of strategies and mechanisms.** Anderson (2004) describes Amazon's strategy as follows:

> This is the power of the Long Tail. The companies at the vanguard of it are showing the way with three big lessons:
> Rule 1: Make everything available
> Rule 2: Cut the price in half. Now lower it.
> Rule 3: Help me find it

There is an interesting contrast with respect to the niche specialty bricks-and-mortar stores we increasingly find in the US market. First, contrary to Amazon, they do not make everything available; in fact, they restrict to a very narrow section of the spectrum. Second, as Proposition 4 suggests, they set higher prices, rather than lower prices. One thing they have in common with Amazon is that they effectively help consumers search, though in a different way.

Interestingly, the drivers for the economic appeal of niche titles are reversed in our work compared to Anderson (2004). In Anderson (2004), it is the lack of capacity constraints that makes it economically viable for large retailers to stock increasingly obscure titles. Conversely, we argue that it is precisely the presence of capacity constraints that motivates small retailers to specialize in narrow niches. Given small stocking capacity, it can be optimal to excel at one niche and neglect all others rather than to be passable at everything.

■ **Bookshop.** Anderson (2004) goes on to argue that

> Most successful businesses on the Internet are about aggregating the Long Tail in one way or another. ... By overcoming the limitations of geography and scale, ... [they] have discovered new markets and expanded existing ones.

One interesting instance of this is given by Bookshop, a relatively recent newcomer in the US book market (Alter, 2020). In essence, Bookshop aggregates local bookstores' catalogues and offers quick, efficient shipping to try and replicate Amazon's business model, while supporting small businesses. Andy Hunter, Bookshop's founder, pitched the e-commerce platform as "the indie alternative to Amazon", and claimed it could represent a "boon for independent stores".

It stands to reason that this type of aggregation is all the more powerful the more specialization (and, thus, heterogeneity) there is among bookstores: if all bookstores were stocking the same bestsellers, Bookshop's business model would fail to replicate even a small fraction of Amazon's variety. Since our analysis provides a rationale for the growth in the number of specialized bookstores (in the US and in recent years), it also provides support for Bookshop's strategy. We test for this in Section 4.

■ **Beyond books.** While our primary focus has been on the book retail market, our analysis, as mentioned in the Introduction, extends to other industries as well. Consider the case of Heatonist, a hot sauce specialist with locations in Manhattan and Brooklyn, New York. Heatonist stocks around 150 different hot sauces, almost always by independent, obscure producers. Popular sauces like Sriracha, which can be found at most US supermarkets, are not offered.

A quick search reveals the extreme extent of Heatonist's specialization: among Heatonist's staff picks, some are entirely absent on Amazon, while less than half have amassed more than 50 Amazon reviews as of March 2022. This is an ever greater degree of specialization than that we model in our paper — in which, for simplicity, we posit that Amazon stocks the whole product space, while brick and mortar stores optimize given capacity.

In the limit, the selection of hot sauces purchased on Amazon can become less niche than those sold offline. While that need not be the case in this or other markets (Heatonist, of course, coexists with several supermarkets only selling a few commercially successful varieties of hot sauces), we show in Section 4 that, in the context of books, this is more than a theoretical possibility.

4. Empirical evidence

Our theoretical results imply a series of predictions. In this section, we discuss empirical evidence from the bookstore industry, specifically evidence from a novel, proprietary data set provided by a major US publisher. The data includes store-title-level wholesale purchases of titles at a monthly frequency. We do *not* observe sales from each channel to consumers. Rather, we assume orders and sales are highly correlated and use the former as a proxy for the latter. High correlation between orders and sales is a natural assumption, particularly in light of the fact that retailers routinely order the same books repeatedly over time, presumably following stockouts. We also have detailed information on approximately 2,800 independent bookstores, including type of store and address, which we have matched to publicly available geographic and demographic data. For non-independent bookstores we have some aggregate information which we describe below.

Bookstore orders can be divided into four different channels:

- **Online D2C**: Sales made to Amazon.

Table 2
Sales distribution by channel

N	Book Chains	Book Stores	Mass Merch.	Online D2C
# titles	43,887	39,267	12,875	47,903
# copies	127,602,337	31,701,747	171,420,650	163,995,077
copies/title	2,907	807	13,314	3,423
% top 100	11.2	11.1	21.4	14.7
% top 1000	39.9	36.1	71.2	45.7
% top 10000	86.6	83.3	99.9	87.4
% niche sales	2.7	4.6	0.6	3.7

- **Bookstores**: Sales made to independent bookstores. At the title and bookstore level.
- **Book chains**: Sales made to bookstore chains such as Barnes & Noble etc. At the title level, aggregated over all book chain stores.
- **Mass Merchandiser**: Sales made through large non-specialty stores such as Target, Walmart etc, as well as airport bookstores. At the title level, aggregated over all book chain stores.

As explained in the above list, our data is at the bookstore level for independent bookstores but aggregate over all bookstores for the other channels. Accordingly, we divide our empirical analysis into two parts. First, we compare aggregate data across the four channels considered above. Second, we dive into store-level data for independent bookstores, testing some of the implications of our theoretical results at the bookstore level.

■ **Aggregate bookstore data.** Proposition 1 predicts that, as Amazon increases in size, bricks-and-mortar stores, especially smaller ones, become increasingly specialized. Extending Proposition 1 to the case of mainstream and niche genres, Proposition 2 implies that bricks-and-mortar sales are more niche-concentrated than online sales (or total sales), despite (or rather, because of) bricks-and-mortar stores' relatively small size.

Table 2 provides a series of indicators that compare different distribution channels and help address the above predictions. The first row shows the number of different titles *collectively* carried by each channel. As can be seen, the value is approximately equal for all channels (about 40 thousand titles) except for mass merchants, who only carry about one third of the titles carried by other channels.

The second row of Table 2 shows the total number of copies sold by each channel over the 2016–2019 period. Dividing this value by the number of copies, we get the average number of copies per title sold by each title. Now we notice that the value for book chains is not that different from that of Amazon. By contrast, independent bookstores sell about third of copies per title of what book chains or Amazon do. By contrast, mass merchants sell about four times more copies per title than book chains or Amazon.

One simple way to test our theory's predictions is to compute concentration indexes by type of channel. To this end, we ask: What does the distribution of sales look like? How does it differ across channels? Specifically, we compute the percentage of sales due to the top N books, where N is equal to 100, 1000, and 10000. As can be seen from Table 2, *in aggregate terms*, book chains are not that different from bookstores in terms of the fraction of sales accounted by the top 100, top 1000, and top 10000 titles. By contrast, mass merchants make twice as many sales from the top 1000 titles. In fact, the distribution of sales by ranking dominates (in the sense of first-order stochastic dominance) that of the other channels.

Finally, and particularly relevant for testing Proposition 2, the bottom row of Table 2 shows that fraction of aggregate sales that corresponds to niche genre titles. We define a niche genre as one with genre market share below the median market share of all genres. Table 3 lists all genre categories and their respective shares of total sales. Comparing the mass merchant and the online columns, we observe numbers that are very consistent with the long-tail narrative, namely that online sales provide an "embarrassment of niches", thus leading to a high fraction of niche titles sold. However, it is equally impressive and relevant that the fraction of niche sales is highest for independent bookstores, namely 4.6%. This is 24% more than Amazon, 70% more than chains, and 660% more than mass merchandisers (which, unsurprisingly, almost exclusively order more familiar titles of more familiar genres).

Unfortunately, we only have aggregated data for the channels book chains and mass merchants. In other words, the data is aggregated over all stores. We do, however, have access to store level data for independent bookstores. This allows us to devise sharper tests of our theoretical results, to which we turn next.

■ **Independent bookstores.** We now focus on sales to independent bookstores, for which we have store-level data. Since purchases are rather sparse (i.e., there are many zeros), we aggregate orders at the title-author level, over time (2016–2019), and across multiple versions of each title (i.e., hardcover or paperback). This results in a sample of 39,000 unique book titles purchased by 2209 unique stores, for a total of 6 million transactions.[8] Table 4 provides some summary statistics of our independent-store-level data. As can be seen, there is significant variation in the extent to which stores carry popular genres (genres with a market share greater than 5%) as well as significant variation in the extent to which stores carry niche genres (genres with a market share lower than .25%).

Figure 5 provides additional characterization of our set of independent stores. It shows the histogram as well as the kernel density of the variable catalog size. Recall that this is based on orders during 2016–2019. Therefore, our catalog size variable is a lower bound of the actual number of titles carried by each store. Considering the large variation in the catalog sizes across stores, we measure this variable in a logarithmic axis. As can be seen from Figure 5, the density is approximately bimodal. Specifically, one can identify two groups of bookstores, one with a catalog of less than 100 titles, one with a catalog of more than 1000 titles.[9]

Our simple theoretical model assumes that each bookstore carries one copy of each title that it carries. This is obviously a simplifying assumption. In reality, one would expect bookstores to carry multiple copies of each title, especially more popular titles. Figure 6 shows the correlation between catalog size and sales (number of copies ordered). As can

8. Each transaction typically includes multiple copies of a given format of a given title on a given date.

9. Again, we note that our measures of size are based on orders during the 2016–2019 period. Therefore, the value of catalog size provides a lower bound of the actual catalog of titles carried by the bookstore.

Table 3
List of book genres and shares of total

Subgenre	Sales (units)	Share of total
Fiction	231,189,777	43.592
Juvenile Fiction	86,342,965	16.280
Biography & Autobiography	33,603,160	6.336
Young Adult Fiction	29,578,303	5.577
Cooking	18,774,375	3.540
Business & Economics	17,566,118	3.312
History	13,219,540	2.493
Juvenile Nonfiction	11,676,213	2.202
Self Help	10,698,177	2.017
Religion	8,984,798	1.694
Health & Fitness	6,018,599	1.135
Political Science	5,622,655	1.060
Social Science	5,542,616	1.045
Family & Relationships	4,030,002	0.760
Science	3,949,072	0.745
Psychology & Psychiatry	3,839,287	0.724
Humor	3,830,617	0.722
Poetry	3,285,440	0.619
True Crime	2,615,300	0.493
Sports & Recreation	2,598,292	0.490
Body, Mind & Spirit	2,475,649	0.467
House & Home	2,392,920	0.451
Young Adult Nonfiction	2,228,676	0.420
Study Aids	2,043,606	0.385
Philosophy	1,476,085	0.278
Art	1,409,606	0.266
Nature - General	1,385,938	0.261
Literary Collections	1,313,664	0.248
Drama	1,242,382	0.234
Travel	1,200,033	0.226
Comics & Graphic Novels	1,154,601	0.218
Language Arts	1,118,371	0.211
Reference	1,005,376	0.190
Performing Arts	703,873	0.133
Law	649,963	0.123
Pets	594,567	0.112
Games	592,391	0.112
Education	586,524	0.111
Crafts & Hobbies	578,725	0.109
Music	539,748	0.102
Medical	518,157	0.098
Photography	421,458	0.079
Computers	394,280	0.074
Technology	353,275	0.067
Mathematics	211,752	0.040
Design	158,232	0.030
Literary Criticism & Collectns	132,800	0.025
Gardening	128,020	0.024
Transportation	123,977	0.023
Antiques & Collectibles	107,778	0.020
Architecture	70,984	0.013
Foreign Language Study	54,450	0.010
Non-Classifiable	15,917	0.003
Bibles	6,036	0.001

Table 4
Summary statistic of independent store level data (2209 stores).

	sales	pop genre sales	niche genre sales	catalog size	% pop gen sales	% niche gen sales
Min	0	0	0	1	0	0
Median	1108	712	21.0	248.5	71.35	1.992
Mean	13226	8646	517.8	5055.5	62.63	4.931
Max	1609174	537703	57316	136116	100.00	100.000

Figure 5
Catalog size distribution

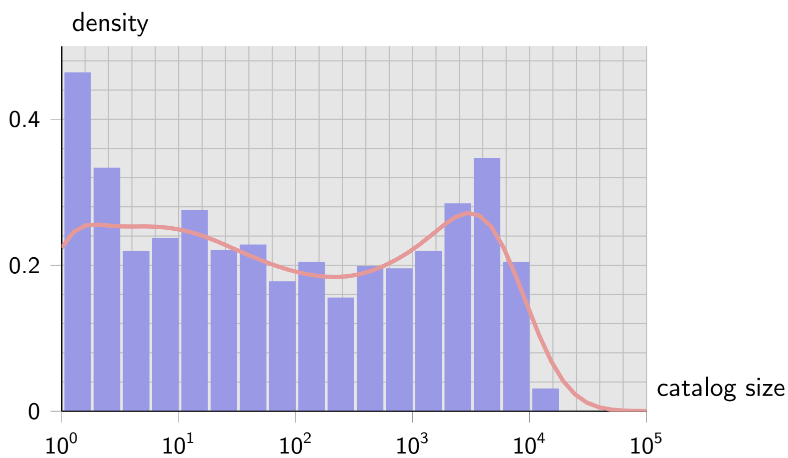

Figure 6
Correlation between catalog size and sales volume

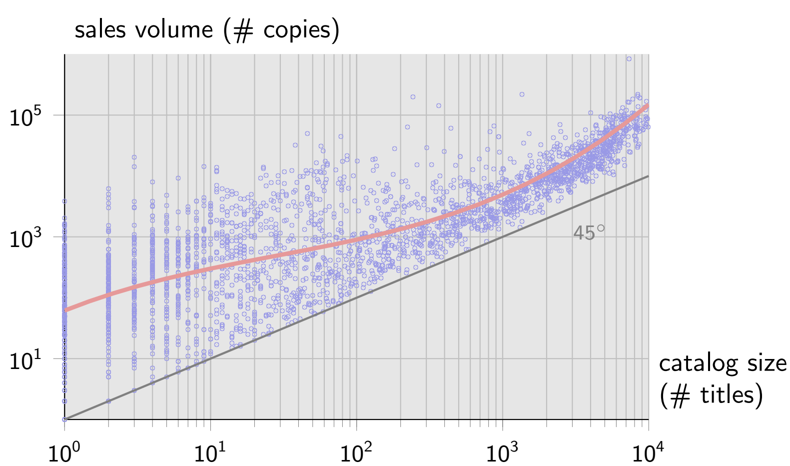

Figure 7
Fraction of niche-genre and popular-genre sales as a function of catalog size

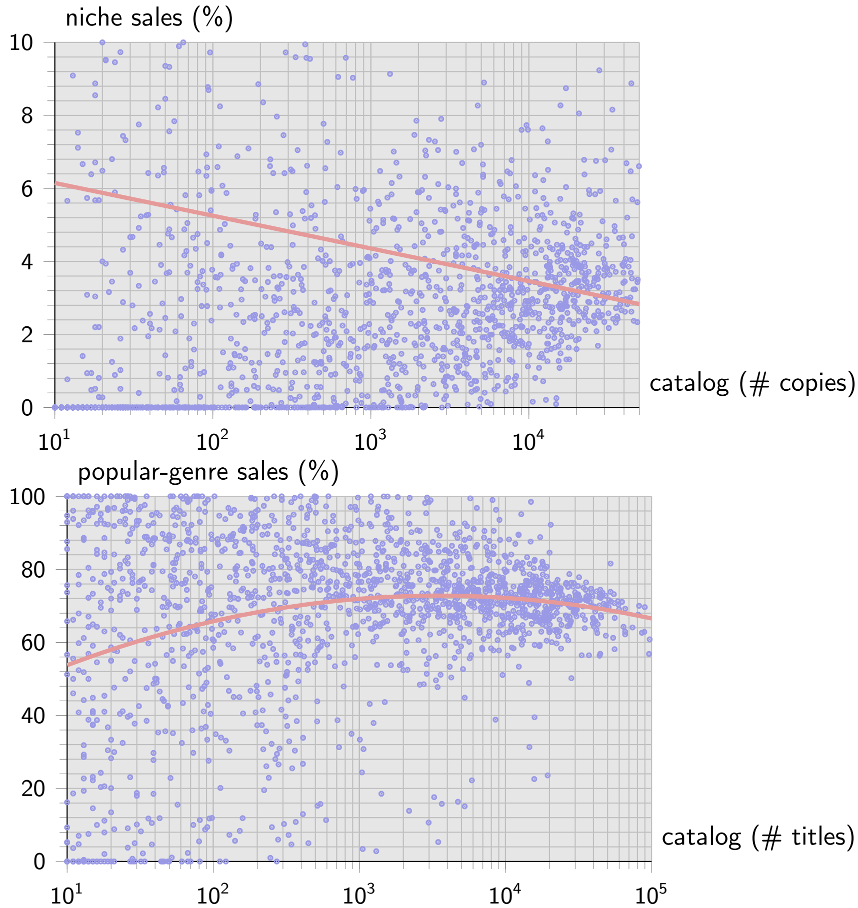

be seen, there is a clear positive correlation between the two measures, which suggests that catalog size and sales are two similar measures of bookstore size. The line in red corresponds to a third-degree polynomial fit, whereas the 45° line provides a lower bound on the relation between the two measures (the number of copies must be greater than the number of titles).

We now get to testing empirical implications from our theoretical results. Broadly speaking, we framed our theory as describing the comparative statics with respect to an increase in z (store a's size in terms of number of titles). However, as Propositions 1 and 2 show, we can think equivalently about the comparative statics with respect to k (store b's capacity in terms of number of titles). This is important because we think of z as varying over time, but our data is limited to a short period (2016–2019). This implies that there is more promise in examining the comparative statics of variation in k than in the evolution of z.

In particular, in the spirit of Proposition 2, we test whether smaller stores are more likely to follow a niche strategy. Figure 7 displays a scatter plot with catalog size (in logs) on the horizontal axis and the seller's concentration on niche titles on the vertical axis. Specifically, we define a niche genre as one with genre market share below the median market share of all genres, which is about .25%. The vertical axis then measures the percentage of the bookstore's sales that correspond to niche-genre sales.

Figure 8
Kernel density of fraction of niche-genre and popular-genre sales

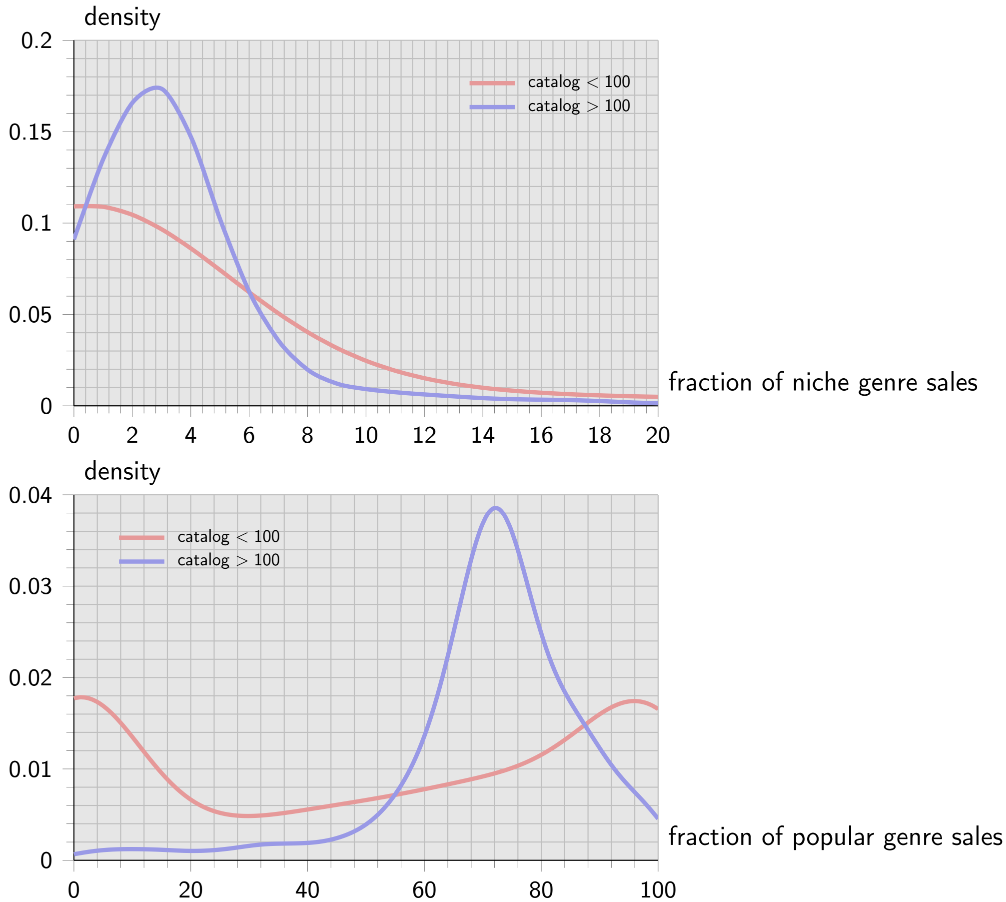

While the scatter plot shows considerable variance, we note a clear negative correlation, as predicted by our theory. A linear regression (on the logarithm of catalog size) shows an economically and statistically significant negative coefficient, as can be seen from the first column in Table 5.

An alternative way of looking at the relation between catalog size and the degree to which bookstores sell niche titles is provided by Figure 8. On the top panel, we see the kernel density of the variable percentage of niche titles for small stores (those with less than 100 titles) and large stores (those with more than 100 titles). The main message of this panel is that the right tail of the distribution corresponding to small stores is considerably thicker than that of larger stores, as predicted by our theory. On the bottom panel, we see the kernel density of the variable percentage of popular-genre titles for small stores (those with less than 100 titles) and large stores (those with more than 100 titles). The main messages from this panel are, first, that the variance of the fraction of popular-genre sales is considerably higher for smaller stores. In fact, the distribution for small stores is bimodal. And second, that the average fraction of popular-genre titles is greater for larger stores, as predicted by Proposition 2.

The regression results presented in Tabletab:regressions suggest that our theory's predic-

Table 5
Explaining the share of niche-genre and popular genre sales

Dependent variable:	niche	niche	popular	popular
log catalog	-.390 (.083)		9.064 (.712)	
log catalog squared			-.555 (.065)	
log sales		-3.322 (.506)		7.739 (1.185)
log sales squared		.196 (.036)		-.336 (.084)
Intercept	7.05 (.528)	17.36 (1.675)	35.84 (1.614)	26.85 (3.927)
Adjusted R-squared	.009	.027	.143	.071
N	2209	2209	2209	2209

Standard errors in parentheses. All coefficient p values lower than .000001

tion is fairly robust. We consider two measures of nich-strategy (choice of niche-genre titles and choice of popular-genre titles) as well as two different measures of store size (catalog size and sales). In all cases, the estimated coefficients are statistically significant and take the sign predicted by theory.

We conclude this section with an important remark regarding the connection between theory and empirical evidence. For the sake of tractability, in our theoretical model all books of a given genre are *ex-ante* equal from a quality perspective. In particular, stores make conscious decisions of which genres to stock but, within a given genre, titles are selected randomly from the available set of titles. This implies that, from an ex-ante point of view, all titles are equally likely to be selected.

Given this simplifying assumption, our model does not, strictly speaking, generate a long tail in sales in the Anderson (2004) sense. The main force behind our offline long tail is the choice by bricks-and-mortar stores of which genres to stock. One can think of a more complex and realistic model in which books (including those of the same genre) differ in terms of ex-ante quality, and where more popular genres have a greater abundance of high-quality titles with potential to become bestsellers. In this context, our results of bricks-and-mortar niche-genre strategy would likely go hand-in-hand with an offline long tail pattern. This would be the more proper sense of the idea of an offline long tail mentioned earlier.

As we saw earlier in this section, our data does provide support both of these complementary interpretations: Not only is the share of sales attributed to niche genre decreasing in store size, but the percentage of sales contributed by the top 10 (top 100, top 1000, and so on) of titles is lower at independent bookstores than it is at Amazon, in spite of the bookstores' limited size. This points to widespread store heterogeneity (if stores were homogeneous, their top 1000 sold titles would account for almost the totality of their sales, given their limited size), which is a consequence of specialization. Moreover — and perhaps even more strikingly — when explicitly plotting a log-sales-rank plot in the spirit of Anderson (2004), we find a slightly longer offline tail than we do online, contrary to conventional wisdom.

5. Conclusion

How can bricks-and-mortar stores survive in an increasingly Amazon-dominated world? In this paper, we suggest that specialization on increasingly narrow niches represents a fundamental strategy to do so. Examples of highly specialized offline retailers abound. For example, Arkipelago in San Francisco exclusively sells Filipino books, while Sweet Pickle Books in the Lower East Side of New York sells pickles and used books, as an homage to the neighborhood's history. Outside of the book industry, we have discussed Heatonist's example – only one of many success stories in boutique food retailing.

Specialization, of course, comes at a steep cost: by specializing in a niche genre that only appeals to a few consumers, bricks-and-mortar stores automatically lose a majority of their potential buyers. However, we show that, as Amazon grows, and particularly for smaller stores, this is a price worth paying: it is better to strongly appeal to some consumers and be ignored by others than to leave all consumers lukewarm. This conclusion is robust to (and, in fact, strengthened by) a variety of extensions, including endogenous prices and offline competition.

Last, our theory allows us to revisit the celebrated long tail theory of Anderson (2004), and to add two novel elements to it: first, while the online long tail has been shown to grow longer and longer over time, we argue that it is unclear whether it is growing *relatively* longer than the offline long tail, contrary to Andersonz's central claim. Second, this implies that Amazon's impact on the rise of niche consumption has been, if anything, understated, as it has neglected Amazon's central role in giving rise to an offline long tail.

Appendix

Proof of Proposition 1: Part (a): Consider the case of a general bookstore. For a x (or y) reader, visiting b yields expected value

$$\tilde{w}+m(k/2)$$

By contrast, buying at a yields expected value

$$m(z/2)$$

given that half of the total titles correspond to genre x (or y). The indifferent buyer is characterized by

$$\tilde{w}=m(z/2)-m(k/2)$$

whenever $m(z/2)-m(k/2)<w$. (Otherwise, every consumer strictly prefers seller a and b makes zero profits.) Finally, b's expected profit (when strictly positive) is given by

$$\pi_g=1-\big(m(z/2)-m(k/2)\big)\,/w \tag{3}$$

Consider now the case of a bookstore specializing in genre x. For an x reader, visiting b yields expected value

$$\tilde{w}+m(k)$$

For a y reader, the value of the x specialty store is zero. As before, buying at a yields expected value

$$m(z/2)$$

both for x and for y readers. The indifferent x buyer is now characterized by

$$\tilde{w}=m(z/2)-m(k)$$

whenever $m(z/2)-m(k)<w$. (Otherwise, every consumer strictly prefers seller a and b makes zero profits.) Finally, b's expected profit (when strictly positive) is given by

$$\pi_s=\tfrac{1}{2}\Big(1-\big(m(z/2)-m(k)\big)\,/w\Big) \tag{4}$$

(Note that, by specializing, b expects to make, at most, $\frac{1}{2}$ in sales. This is because it will have lost all potential readers from the genre it did not specialize in.)

If $z=0$, that is, if Amazon is out of the picture, then being a general store is trivially a dominant strategy: the store sells to a measure 1 of consumers, whereas the specialty store sells to a measure $\frac{1}{2}$ only (at the same price). Specifically, a general store's profits are equal to 1, the highest value possible, while a specialty store would only achieve its upper bound, $\frac{1}{2}$.

At the opposite end, let z_g is such that $\left(m(z_g/2)-m(k/2)\right)/w=1$. For $z=z_g$, we have $\pi_g=0$, whereas

$$\pi_s=\tfrac{1}{2}\,\big(1-\big(m(z_g/2)-m(k)\big)\,/w\big)>\tfrac{1}{2}\,\big(1-\big(m(z_g/2)-m(k/2)\big)\,/w\big)=0$$

Such an z will exist whenever $\lim_{z\to\infty}\big(m(z/2)-m(k/2)\big)\,/w>1$, which is implied by the condition in the Proposition. (As mentioned in the text, if this condition does not hold —

for instance because w or k are very large, or $m(n)$ is very flat —, then it may always be optimal for the store to be generalist.)

Given continuity of π_g and π_s, it follows from the intermediate value theorem that there exists an $z_{gs} \in (0, z_g)$ such that $\pi_g(z_{gs}) = \pi_s(z_{gs})$, where for notational simplicity we have suppressed the store profit's dependence on k and w. To show that z_{gs} is unique we note that

$$\frac{d\,(\pi_s - \pi_g)}{d\,z} = \big(- m'(z/2) + 2\,m'(z/2)\big)\,/(4\,w) = m'(z/2)/(4\,w) > 0 \tag{5}$$

where the inequality follows from the fact that $m(z)$ is strictly increasing for every z. This concludes the first part of the proof.

To show that $z_{gs}(k,w)$ increases in k and w, we compute the derivative of the profit difference $(\pi_s - \pi_g)$ with respect to k and w:

$$\frac{\partial\,(\pi_s - \pi_g)}{\partial\,k} = \frac{m'(k)}{2\,w} - \frac{m'(k/2)}{2\,w} = \frac{1}{2\,w}\,\big(m'(k) - m'(k/2)\big) < 0 \tag{6}$$

where the inequality follows from concavity of m (David, 1997). Similarly,

$$\frac{\partial\,(\pi_s - \pi_g)}{\partial\,w} = \frac{m(z/2) - m(k)}{2\,w^2} - \frac{m(z/2) - m(k/2)}{w^2} = (\tfrac{1}{2} - \pi_s)/w - (1 - \pi_g)/w$$

where the second equality follows from (3) and (4). By definition, $\pi_s = \pi_g = \overline{\pi}$ at $z = z_{gs}$. It follows that

$$\left.\frac{\partial\,(\pi_s - \pi_g)}{\partial\,w}\right|_{z\,=\,z_{gs}} = (\tfrac{1}{2} - \overline{\pi})/w - (1 - \overline{\pi})/w = -1/(2\,w) < 0 \tag{7}$$

By the implicit function theorem,

$$\frac{\partial\,z_{gs}(k,w)}{\partial\,k} = -\frac{\partial\,(\pi_s - \pi_g)\,/\,\partial\,k}{\partial\,(\pi_s - \pi_g)\,/\,\partial\,z} > 0$$

where the inequality follows from (5) and (6). Also by the implicit function theorem,

$$\left.\frac{\partial\,z_{gs}(k,w)}{\partial\,w}\right|_{z\,=\,z_{gs}} = -\frac{\partial\,(\pi_s - \pi_g)\,/\,\partial\,w\,\big|_{s\,=\,z_{gs}}}{\partial\,(\pi_s - \pi_g)\,/\,\partial\,z} > 0$$

where the inequality follows from (5) and (7).

Part (b): We have that

$$\frac{\partial(\pi_g - \pi_s)}{\partial k} = \frac{1}{2}m'(k/2) - \frac{1}{2}m'(k) > 0$$

by concavity of k. Moreover, we know that, as $k \to z$, $\pi_g \to 1$, $\pi_s \to 1/2$, and thus $k_g > k_s$ whenever k is large enough.

Conversely, we know that $\pi_g = 0$ whenever $m(z/2) - m(k/2) \geq w$, while $\pi_s = 0$ whenever $m(z/2) - m(k) \geq w$. Denote by k_g^* and k_s^* the two values of k that satisfy these two with equality. Because both expressions are decreasing in k, these exist and are non-negative if and only if $m(z/2) \geq w$, which is implied by the condition in the proposition.

Now, notice that $k_g^* = 2k_s^*$. Thus, whenever k_g^* and k_s^* are positive, we have that $k_g^* > k_s^*$ or, in other words,

$$\pi_s > \pi_g = 0, \quad \forall k \in [k_s^*,\ k_g^*].$$

Combining our observations, we have that the difference $\pi_g - \pi_s$ is negative for $k \in [k_s^*,\ k_s^*]$ and monotonically increases, becoming strictly positive for $k \to s$. Thus, there exists a unique k_{gs} such that $\pi_s(k_{gs}, s) = \pi_g(k_{gs}, s)$.

Now, we want to show that $k_g s(z, w)$ is decreasing in z and increasing in w. To do so, we appeal to the Implicit Function Theorem again, in a similar fashion as in part (a).

We have that

$$\frac{\partial k_{gs}(z,w)}{\partial s} = -\frac{\partial(\pi_g - \pi_s)/\partial s}{\partial(\pi_g - \pi_s)/\partial k} > 0$$

and

$$\frac{\partial k_{gs}(z,w)}{\partial w} = -\frac{\partial(\pi_g - \pi_s)/\partial s}{\partial(\pi_g - \pi_s)/\partial w} < 0,$$

which concludes part (b) of the proof. ■

Proof of Proposition 2: Suppose store b specializes in genre x, the popular genre ($\alpha > \frac{1}{2}$). Then store b reaches at most α of its potential customers. The indifferent customer (indifferent between store a and store b) has z such that

$$m(\alpha\, z) = m(k) + \tilde{w}$$

where $\alpha\, z$ is total supply of titles of genre x, all of which are available at store a; and k is the supply of titles of genre x at store b (in other words, all of store b's capacity, k, is devoted to carrying genre x titles). It follows that, of the k store-b potential customers, a fraction $\alpha\, k$ is interested in the genre offered by store b, and a fraction $\big(m(\alpha\, z) - m(k)\big)/\tilde{w}$ of this fraction prefers store b to store a. This implies that store b's profit from specializing in genre x is given by

$$\pi_x = \alpha\Big(1 - \big(m(\alpha\, z) - m(k)\big)/\tilde{w}\Big)$$

Similarly, the profit from specializing in genre y is given by

$$\pi_y = (1-\alpha)\Bigg(1 - \Big(m\big((1-\alpha)\, z\big) - m(k)\Big)/\tilde{w}\Bigg)$$

If $z = 0$, that is, if Amazon is out of the picture, then the popular genre x is trivially a dominant strategy: the store sells to a measure α of consumers, whereas the niche-genre store sells to a measure $1-\alpha < \alpha$ only (and at the same price). At the opposite end, let z_x be the value of z such that $\pi_x = 0$. Such an z will exist whenever $\lim_{z\to\infty}\big(m(\alpha\, z) - m(k)\big)/w > 1$, which is equivalent to the condition in the Proposition. We then have

$$\pi_y = (1-\alpha)\Bigg(1 - \Big(m\big((1-\alpha)\, z_x\big) - m(k)\Big)/\tilde{w}\Bigg) > \alpha\Big(1 - \big(m(\alpha\, z_x) - m(k)\big)/\tilde{w}\Big) = 0$$

(If this condition does not hold — for instance because w or k are very large, or $m(n)$ is very flat —, then it may always be optimal for the store to choose the popular genre.)

Given continuity of π_x and π_y, the intermediate value theorem implies that there exists at least one value $\widehat{z}_{xy} \in (0, z_x)$ such that $\pi_g(\widehat{z}_{xy}) = \pi_s(\widehat{z}_{xy})$, where for notational simplicity we have suppressed the store profit's dependence on k and $\tilde{w}$. Let z_{xy} be the highest of these values. Then $\pi_y \geq \pi_x$ for $z > z_{xy}$.

Consider now the comparative statics with respect to k. First notice that there are only $(1-\alpha)\,z$ titles of genre y. Therefore, $m\big((1-\alpha)\,z_x\big)$ is an upper bound of the benefit from stocking only y genre titles. Therefore, for $k > (1-\alpha)\,z_x$, $\pi_y = (1-\alpha)\,k$. As to π_x, we can see that it is increasing in k and, as k reaches $k = m(\alpha\,z)$, $\pi_x = \alpha\,k > \pi_y$. It follows that there exist a k_{xy} such that $\pi_x > \pi_y$ if $k > k_{xy}$. ■

Proof of Proposition 3: From Proposition 1, there exists a threshold $k_{gs}(z,w)$ such that an active firm b optimally chooses to be a specialty store if and only if $k < k_{gs}(z,w)$. Define

$$k_1 = k_{gs}(z',w)$$
$$k_2 = k_{gs}(z'',w)$$

From Proposition 1, $k_{gs}(z,w)$ is increasing in z, implying that $k_2 > k_1 > 0$.

From (1), concavity of $m(k)$ implies concavity of $\pi_g(z,k)$ with respect to k. Moreover, Assumption **??** implies that $\lim_{k\to 0}\ \pi_g(z,k) > 0$. Finally, since revenues are bound but costs are not, $\lim_{k\to\infty}\ \pi_g(z,k) = -\infty$. Together, these properties allow us to define k_3 and k_4 as

$$\pi_g(z'',k_3) = 0$$
$$\pi_g(z',k_4) = 0$$

We can establish that $\pi_g(z'',k) > 0$ if and only if $k < k_3$ and $\pi_g(z',k) > 0$ if and only if $k < k_4$. Condition (2) implies that $\pi_g(z'',k_2) > 0$, which in turn implies that $k_2 < k_3$. Finally, monotonicity of π_g with respect to z implies that $k_4 > k_3$. We can thus establish that the values of k as defined above satisfy $0 < k_1 < k_2 < k_3 < k_4$.

If $0 < k < k_1$, then Proposition 1 implies that a store strictly prefers to be a specialist for both $z = z'$ and $z = z''$. This is because $k < k_1 = k_{gs}(z',w)$ (implying a specialist choice for $z = z'$) and $k < k_2 = k_{gs}(z'',w)$ (implying a specialist choice for $z = z''$), where $k_2 > k_1$. Moreover, since $k_1 < k_3$ and $k_1 < k_4$, a generalist store makes positive profits when $0 < k < k_1$. Since specialist profits are greater than generalist profits, we conclude that a specialist store makes positive profits. Together, the above properties imply that, if $0 < k < k_1$, then it is optimal to be an active specialist store for $z = z'$ and $z = z''$.

If $k_1 < k < k_2$, then at z' the store is better off as a generalist (because if $k > k_1$), while at z'' it is better off as a specialist (because $k < k_2$). Since $k_2 < k_3$ and $k_2 < k_4$, both types of store prefer to be active than to exit.

If $k_2 < k < k_3$, then both an z' store and an z'' store prefer to be a generalist (because $k > k_2$). Moreover, both z' and z'' prefer to be active (because $k < k_3$ and $k < k_4$).

Finally, if $k_3 < k < k_4$, then both an z' store and an z'' store prefer to be a generalist (because $k > k_2$). An z' store prefers to be active (because $k < k_4$) but an z'' store prefers to be inactive (because $k > k_3$). ■

Proof of Proposition 4: We first solve for the optimal prices of a general store given that store a sets p_a. Store g's profit is given by $\pi_g = p_g\,q_g$, where q_g, the store's sales, are given by

$$q_g = 1 - \big(m(z/2) - m(k/2) - p_a + p_g\big)\,/w$$

The profit-maximizing price, quantity and profit levels are given by

$$\widehat{p}_g = \tfrac{1}{2}\big(w - m(z/2) + m(k/2) + p_a\big) \tag{8}$$

$$\widehat{q}_g = \tfrac{1}{2}\big(w - m(z/2) + m(k/2) + p_a\big)/w = \widehat{p}_g/w \tag{9}$$

$$\widehat{\pi}_g = \widehat{p}_g\,\widehat{q}_g = (\widehat{p}_g)^2/w \tag{10}$$

In the case of a specialty store, profit is given by $\pi_s = p_s\,q_s$, where q_s, the store's sales, are given by

$$q_s = \tfrac{1}{2}\Big(1 - \big(m(z/2) - m(k) - p_a + p_s\big)/w\Big)$$

The profit-maximizing price, quantity and profit levels are given by

$$\widehat{p}_s = \tfrac{1}{2}\big(w - m(z/2) + m(k) + p_a\big) \tag{11}$$

$$\widehat{q}_s = \tfrac{1}{4}\big(w - m(z/2) + m(k) + p_a\big)/w = \widehat{p}_s/(2\,w) \tag{12}$$

$$\widehat{\pi}_s = \widehat{p}_s\,\widehat{q}_s = (\widehat{p}_s)^2/(2\,w) \tag{13}$$

Direct inspection of (8) and (11) reveals that

$$\widehat{p}_s > \widehat{p}_g$$

that is, in equilibrium specialty bookstores set a higher price. Moreover, from (8)–(9) and (11)–(12) we conclude that

$$\widehat{p}_s/\widehat{q}_s = 2\,w > \widehat{p}_g/\widehat{q}_g = w \tag{14}$$

Consider the extreme case when $z = 0$. Straightforward computation shows that $\widehat{\pi}_g > \widehat{\pi}_s$ if and only if the condition in the Proposition holds. At the opposite end, let z_g be such that $\widehat{p}_g = 0$. Comparing (8) and (11), we see that, at $z = z_g$, $\widehat{p}_s > \widehat{p}_g = 0$. From (10) and (13) we conclude that, at $z = z_g$, $\widehat{\pi}_s > \widehat{\pi}_g = 0$. Since both $\widehat{\pi}_s$ and $\widehat{\pi}_g$ are continuous we conclude by the intermediate-value theorem that there exists at least one $\widetilde{z}_{gs}$ such that $\widehat{\pi}_s = \widehat{\pi}_g$. Let z_{gs} be the highest of these values. Then $\widehat{\pi}_s > \widehat{\pi}_g$ when $z_{gs} < s < z_g$.

Finally, notice that, at $z = z_{gs}$, $\widehat{\pi}_g = \widehat{\pi}_s$, that is, $\widehat{p}_g\,\widehat{q}_g = \widehat{p}_s\,\widehat{q}_s$. Since, from (14), $\widehat{p}_s/\widehat{q}_s > \widehat{p}_g/\widehat{q}_g$, it must be that, at $z = z_{gs}$, $\widehat{p}_s > \widehat{p}_g$ and $\widehat{q}_s < \widehat{q}_g$. Since these are strict inequalities, they also hold in the neighborhood of $z = z_{gs}$. It follows that, in the right neighborhood of $z = z_{gs}$, a specialty store earns a higher profit, sets a higher price, and captures a lower market share. ■

Proof of Proposition 5: The proof follows straightforwardly from the definitions of π_g and π_s. Specifically, from (3) and (4) we derive

$$\begin{aligned}\frac{\partial \pi_g}{\partial w} &= \frac{m(z/2) - m(k/2)}{w^2}\\ \frac{\partial \pi_s}{\partial w} &= \frac{m(z/2) - m(k)}{2\,w^2}\end{aligned} \tag{15}$$

which implies that

$$\frac{\partial \pi_s}{\partial w} = \tfrac{1}{2}\,\frac{\partial \pi_s}{\partial w} > 0$$

Taking derivatives of (15) with respect to z, we get

$$0 < \frac{\partial^2 \pi_s}{\partial w\, \partial z} = \tfrac{1}{2}\, \frac{\partial^2 \pi_s}{\partial w\, \partial z}$$

Taking derivatives of (15) with respect to k, we get

$$0 < \frac{\partial^2 \pi_s}{\partial w\, \partial k} = \tfrac{1}{2}\, \frac{\partial^2 \pi_s}{\partial w\, \partial k}$$

Finally, taking derivatives of (15) with respect to w we get

$$\frac{d^2 \pi_g}{d w^2} < \frac{d^2 \pi_s}{d w^2} < 0$$

which concludes the proof. ■

Proof of Proposition 6: Figure 9 illustrates the competition case. On the horizontal axis we measure the consumer location d, where $d = 0$ corresponds to bricks-and-mortar store b_0 and $d = 1$ corresponds to bricks-and-mortar store b_1. On the vertical axis we measure z, the relative preference for a bricks-and-mortar store. We assume that d and z are independently and uniformly distributed: $\tilde{d} \sim U[0, 1]$ and $\tilde{w} \sim U[0, w]$. Since there are two different genres, we need to plot one graph per genre, genre x on the top panel and genre y on the bottom panel.

Figure 9 illustrates the case when both b_0 and b_1 are general stores. Store b_0's demand of genre x is given by the area in blue in the top panel, whereas store b_0's demand of genre y is given by the area in red in the top panel. To understand that, notice that store b_0 must beat both store a and store b_1. Beating store a requires

$$m(k/2) + z - \tau\, \tilde{d} > m(z/2)$$

whereas beating store b_1 requires

$$m(k/2) + z - \tau\, \tilde{d} > m(k/2) + z - \tau\, (1 - \tilde{d})$$

This results in the following set of inequalities

$$\begin{array}{rcl} \tilde{w} & > & m(z/2) - m(k/2) + \tau\, \tilde{d} \\ \tilde{d} & < & \tfrac{1}{2} \end{array}$$

which in turn correspond to the areas in blue (top panel) and red (bottom panel).

Given that b_1 chooses to be a general store, how does b_0 change its profits by specializing in genre x? Store b_1's demand from x consumers is now determined by

$$m(k) + \tilde{w} - \tau\, \tilde{d} > m(z/2)$$

(beat firm a) and

$$m(k) + \tilde{w} - \tau\, \tilde{d} > m(k/2) + \tilde{w} - \tau\, (1 - \tilde{d})$$

Figure 9
Store strategy under bricks-and-mortar competition

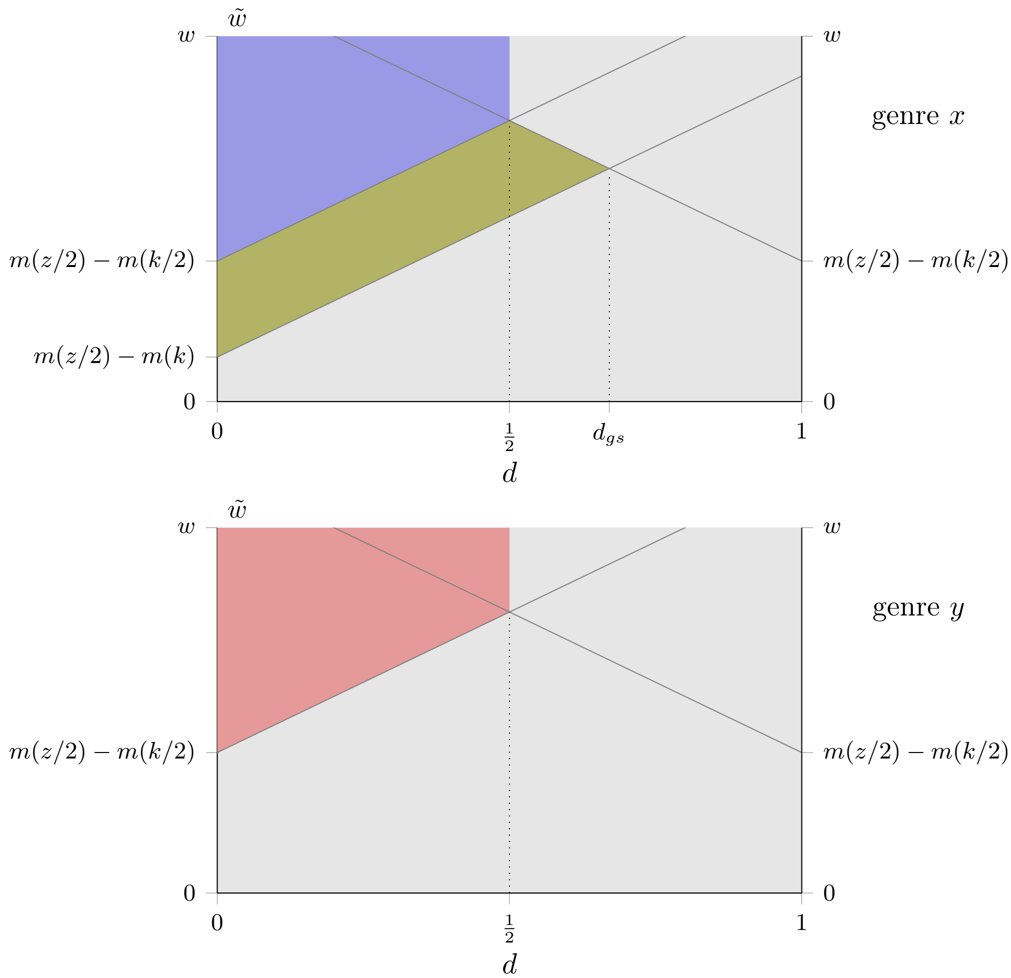

(beat firm b_1). This simplifies to

$$\begin{aligned} \tilde{w} &> m(z/2) - m(k) + \tau\, \tilde{d} \\ \tilde{d} &< d_{gs} \equiv \tfrac{1}{2} + \big(m(k) - m(k/2)\big)\, / \tau \end{aligned}$$

This corresponds to an increase in demand for genre x given by the area in green on the top panel and a loss in demand for genre y given by the area in red on the bottom panel. The green area on the top panel corresponds entirely to consumers who purchased from a when both b_0 and b_1 were general stores and now prefer to buy from b_0, the genre x specialty store. The red area on the bottom panel corresponds to consumers who were interested in store b_0 when it was a general store but are now not interested since it no longer carries any genre y titles.

The values of z and k in Figure 9 were chosen so that the areas in green and red are equal. This implies that, given that store b_1 follows a general-store strategy, store b_0 is indifferent between being a general store and being a specialty store. Suppose now that b_1

Figure 10
Store strategy under bricks-and-mortar competition

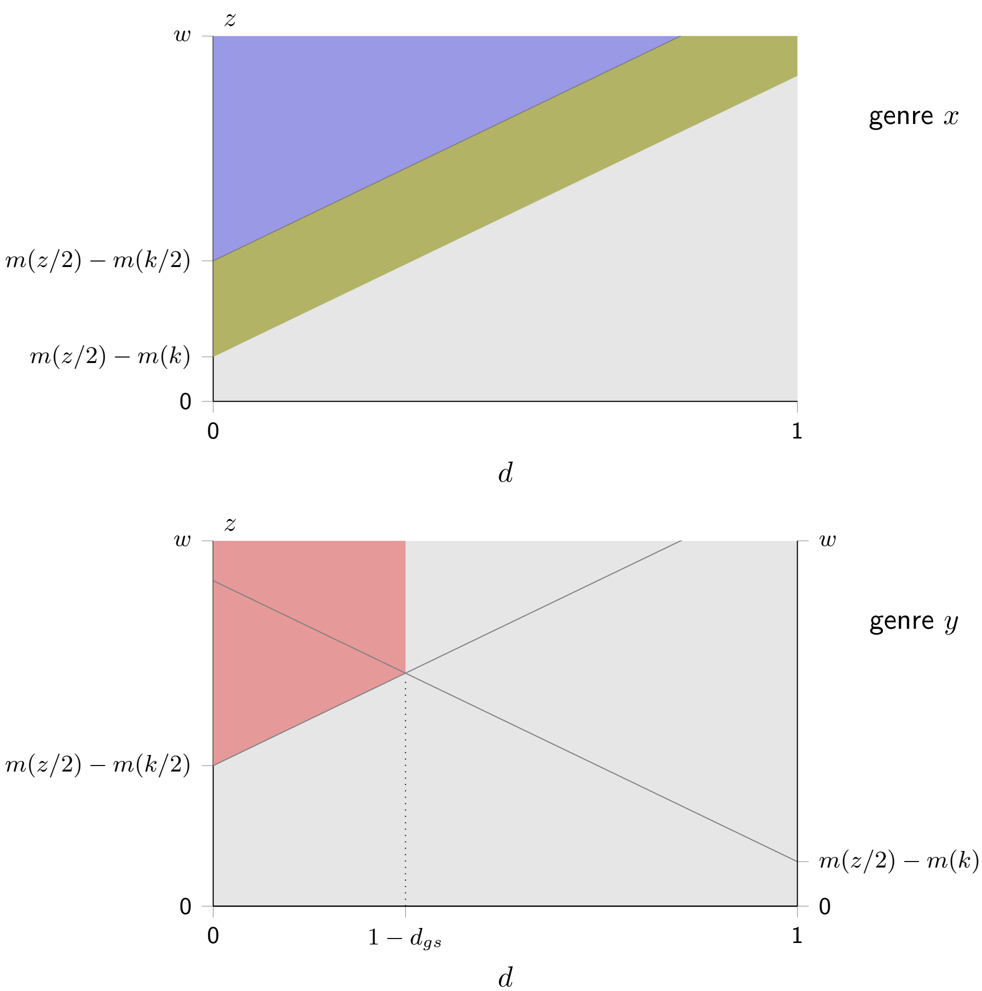

chooses to be a y-specialty store. What is the gain for store b_0 from specializing in x? This alternative scenario is described in Figure 10. In terms of x consumers, the battle is now limited to firms b_0 and a, since firm b_1 is absent from this genre. Demand for firm b_0 is determined by

$$m(k/2) + \tilde{w} - \tau\,\tilde{d} > m(z/2)$$

which corresponds to the area in blue. Regarding genre y (bottom panel), we still need to consider both competition by a and competition by b_1. Since b_1 is a genre y specialty store, we now have

$$\begin{aligned} \tilde{w} &> m(z/2) - m(k) + \tau\,\tilde{d} \\ \tilde{d} &< 1 - d_{gs} \equiv \tfrac{1}{2} + \big(m(k/2) - m(k)\big)\,/\tau \end{aligned}$$

which corresponds to the area in red. What happens to firm b_0's profit as it switches from a general store to a genre x specialty store? On the top panel (that is, in terms of x sales), it experiences a profit increase given by the green area. On the top panel (that is, in terms of y sales), it experiences a profit loss given by the red area.

Immediate inspection reveals that the green area in the top panel of Figure 10 is greater than the green area in the top panel of Figure 9, whereas the red area in the bottom panel of Figure 10 is lower than the red area in the bottom panel of Figure 9. This implies that, if firm b_0 is indifferent between being a general store and being a specialty store when its rival is a general store, then it strictly prefers to be a specialized store when its rival is a specialty store. ■

References

Alter, Alexandra (2020), "Bookstores Are Struggling. Is a New e-Commerce Site the Answer?" *New York Times*, June 16.

Anderson, Chris (2004), "The Long Tail," *Wired Magazine*, October 1.

Bar-Isaac, Heski, Guillermo Caruana, and Vicente Cuñat (2012), "Search, Design, and Market Structure," *American Economic Review*, 102, 1140–60.

Brynjolfsson, Erik, Yu (Jeffrey) Hu, and Duncan Simester (2011), "Goodbye Pareto Principle, Hello Long Tail: The Effect of Search Costs on the Concentration of Product Sales," *Management Science*, 57, 1373–1386.

Chaudhuri, Saabira (2019), "New CEO Wants to Make Barnes & Noble Your Local Bookstore," *Wall Street Journal*, August 8.

Choi, Jeonghye and David R Bell (2011), "Preference Minorities and the Internet," *Journal of Marketing Research*, 48, 670–682.

David, H. A. (1997), "Augmented Order Statistics and the Biasing Effect of Outliers," *Statistics & Probability Letters*, 36, 199–204.

Forman, Chris, Anindya Ghose, and Avi Goldfarb (2009), "Competition between Local and Electronic Markets: How the Benefit of Buying Online Depends on Where You Live," *Management Science*, 55, 47–57.

Goldmanis, Maris, Ali Hortaçsu, Chad Syverson, and Önsel Emre (2010), "E-Commerce and the Market Structure of Retail Industries," *The Economic Journal*, 120, 651–682.

Igami, Mitsuru (2011), "Does Big Drive out Small?" *Review of Industrial Organization*, 38, 1–21.

Johnson, Justin P. and David P. Myatt (2006), "On the Simple Economics of Advertising, Marketing, and Product Design," *American Economic Review*, 96, 756–784.

Kahn, Barbara and Ray Wimer (2019), "What's the next Chapter for Barnes & Noble?" Knowledge@Wharton Podcast, University of Pennsylvania.

Mankiw, N. Gregory and Michael D. Whinston (1986), "Free Entry and Social Inefficiency," *The RAND Journal of Economics*, 17, 48–58.

Neiman, Brent and Joseph S. Vavra (2019), "The Rise of Niche Consumption," Technical Report w26134, National Bureau of Economic Research.

Raffaelli, Ryan (2020), "Reinventing Retail: The Novel Resurgence of Independent Bookstores," HBS Working Paper 20-068, Harvard University.

Rhodes, Andrew and Jidong Zhou (2019), "Consumer Search and Retail Market Structure," *Management Science*, 65, 2607–2623.

Saxena, Jaya (2022), "The Coolest Place to Drink Is Your Local Bookstore," *Eater NY*.

Segal, David (2019), “Can Britain’s Top Bookseller Save Barnes & Noble?” *The New York Times*, April 29.

Todd, Sarah (2019), “Barnes & Noble’s Fate Rests in the Hands of a British Indie Bookstore Owner,” *Quartz*, July 21.

Waldfogel, Joel (2007), *The Tyranny of the Market: Why You Can’t Always Get What You Want*, Harvard University Press, Cambridge, Mass.