Amazon and the Evolution of Retail

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 an offline long tail that is thicker than the online long tail, contrary to existing research. Offline specialization benefits consumers; in fact, consumers would benefit from more specialization than it results in equilibrium. We discuss this and other relevant comparative statics based on a simple model of consumer demand and retail design. We complement our theoretical analysis with corroborative empirical evidence. To do so, we employ 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.

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 advanced data and 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 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 are of 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 the channel they purchase from. 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, both stores' intensive margins 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 player co-existing with multiple niche players and a declining number of mid-size and large bricks-and-mortar stores such as Barnes & Noble [see, e.g., Kahn and Wimer, 2019].

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

> 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 an incredibly high number of increasingly obscure titles. A bookstore that can only store – say – 1000 books, according to

Anderson, will instead opt for 1000 popular, mainstream titles. After all, why use precious and scarce shelf space on books that only attract few potential buyers?

What's missing from this observation and prediction is the endogenously determined bricks-and-mortar store strategy, both in terms of size and – especially – specialization. So, while it is true that an increasing percentage of total sales originate in niche products, our analysis suggests that this is not particularly true for online sellers; in fact it could be particularly true for 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 one, which is the byproduct of offline stores' specialization – itself a counter to Amazon's increasing dominance.

We provide some empirical evidence for our theoretical claims, including in particular a dataset from a large publisher from 2016–2019. By observing all sales made by the publisher to different type of book retailers (independent bookstores, book chains, online retailers, airport bookstores) over this period – for a total of nearly 6 million transactions – we confirm that bricks-and-mortar bookstores have become smaller and more specialized than their competitors, to an extent that, overall, their long tail is longer than Amazon's.

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

■ **Related literature.** Conceptually, the paper that is closest to us 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). As well, 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 a possible 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 look at 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 3 *C's*: 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 interests*") and community. All of these strongly resonate with both our theoretical and empirical findings.

2. Theory

Consider an economy with two book sellers, a (Amazon) and b (bricks-and-mortar); and two different book genres, x and y. 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 v is generated from a cdf $F(v)$, where $f(v) > 0$ if and only if $v \in [0, \overline{v}]$, where $\overline{v}$ is possibly infinite.[2,3]

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[4], a desire to support small and local businesses, or an ideological aversion to (or taste for) Amazon. We assume that this preference is uniformly distributed in $[0, \overline{z}]$.[5]

---

1. Later in the paper, we consider the asymmetric case, that is, the case of a popular genre and a niche genre.
2. For instance, $F(\cdot)$ could be a bounded distribution — like a uniform, which we will use for some of our numerical results — or an unbounded one — like an exponential — reflecting the (rare) presence of arbitrarily good books in each genre.
3. Later in the paper we consider a generalization whereby some consumers have positive valuation for both genres. The qualitative nature of our main results does not depend on our assumption (for much of the paper) that there are no such "eclectic" buyers.
4. Saxena [2022] describes recent examples of independent bookstores providing offline perks such as bars and cafes.
5. The assumption that the lower bound of z 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 z, corresponding to a relative preference for firm a.

Seller a carries all titles in the economy, a total of s titles, $s/2$ 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 v of a title at no cost. By contrast, when b chooses what books to carry, it can observe genre but not v. Therefore, the bookstore determines which type of books to sell but otherwise selects a random sample of values 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 v. If the store carries x titles of the buyer's preferred genre, then the buyer receives an expected value $m(x)$, where $m(x)$ is the expected value of the highest element of a sample of size x drawn from $F(v)$.

■ **General or specialty store?** The focus of our analysis is on bookstore b's strategy as the value of s 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 sells up to $k/2$ titles of each genre, whereas a specialty store can sell up to 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 s (a measure of the online store's growth), and k (size heterogeneity across bricks-and-mortar stores). Our first two results are based on the following assumption:

Assumption 1. *[Interior solution]* $\quad \overline{v} - \overline{z} > m(k)$.

This assumption ensures that the solution in interior.[6] Specifically, when Assumption 1 fails to hold, then we are in a corner solution whereby it is a dominant strategy for b to be a general store. If Assumption 1 holds, however, then the choice of general or specialty store depends on the relative value of s and k, as stated in the following result:

Proposition 1. *[Threshold strategy] Suppose Assumption 1 holds. (a) There exists a threshold $s_{gs} = s_{gs}(k, \overline{z})$ such that an active firm b optimally chooses to be a specialty store if and only if $s > s_{gs}$. Moreover, $s_{gs}(k, \overline{z})$ is increasing in both k and $\overline{z}$. Equivalently, (b) There exists a threshold $k_{gs} = k_{gs}(s, \overline{z})$ such that an active firm b optimally chooses to be a specialty store if and only if $k < k_{gs}$. Moreover, $k_{gs}(s, \overline{z})$ is decreasing in s and increasing in $\overline{z}$.*

Proof: 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 the specialty genre, the store increases the expected quality that a patron expects from visiting the store (intensive margin). As total supply s increases, the expected payoff from visiting store a, $m(s)$, increases. This implies that 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 s past which a specialty store strategy yields higher profit.

---

6. We note that Assumption 1 is trivially satisfied when $\overline{v} = \infty$.

Another way of understanding Proposition 1 is that, as s 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. 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.

Other than s, we also consider comparative statics with respect to k and $\overline{z}$ First, for a given value of s, a store with larger capacity 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.[7] Similarly, the threat posed by Amazon is lower the greater $\overline{z}$, that is, the greater the buyers' aversion to purchasing from Amazon. Accordingly, store b is less likely to become a specialty store as a strategy to cope with online competition.

■ **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 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 αs of the total titles are of genre x, and a fraction $(1 - \alpha) s$ are of genre y. Proposition 1 states that, as s 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 s is high enough.

Proposition 2. *[Niche strategy] Suppose Assumption 1 holds. There exists an s_{xy} such that an active store b specializes in a niche genre (rather than a popular genre) if $s > s_{xy}$.*

Figure 1 illustrates Proposition 2. The key insight is that, *relatively* speaking, a niche-genre store suffers less from an increase in s 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 s, 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 s, however, the niche strategy becomes increasingly attractive, as illustrated by Figure 1. Specifically, for $s > s_{xy}$, 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 s_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 $s > s_{gs}$ is a necessary and sufficient condition

---

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

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

for specialization. By contrast, in Proposition 2 $s > s_{xy}$ is only a 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 $s > s_{xy}$ becomes a necessary and sufficient condition.[8]

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 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 $m(x) = x/(1 + x)$. 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 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 a niche genre (y).

---

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

Fig. 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.

As can be seen, as s 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. Moreover, in order to make reasonable comparative statics with respect to k, we now assume that the measure of consumers who must decide between buying from a or buying from b is equal to k. In other words, we assume the bookstore's technology is characterized by constant returns to scale: both capacity costs and potential consumer reach vary linearly with k. We will return to this assumption later.

Now that we assume $c > 0$, a third option — exit — becomes non-trivial. We consider the bookstore's optimal choice in the (s, c) 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 s_{gs} such that, conditional on being active, a specialty-store strategy is better than a general-store strategy if and only if $s > s_{gs}$. We now consider two additional comparisons: general vs exit and specialty vs exit.

Proposition 3. [exit] *For a given s, firm b's optimal choice is to exit if and only if $c > c^\circ(s)$. Conversely, for a given $c > \underline{c}$, where $\underline{c} \equiv \left(m(2\,k) + \overline{z} - \overline{v}\right)/(2\,\overline{z})$, firm b's optimal choice is to exit if and only if $s > s^\circ(c)$. Finally, if $c < \underline{c}$ then exit never takes place.*

Figure 3 illustrates Proposition 3 in the linear case, that is, $F(v) = v/\overline{v}$. 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 s_{gs} threshold, these lines define three regions: the GENERAL region, defined by $s < s_{gs}$ and $c < c^\circ$ (effectively, the generalist exit boundary); the SPECIALTY region, defined by $s > s_{gs}$ and $c < c^\circ$ (effectively, the specialty exit boundary); and the EXIT region, defined by $c > c^\circ$.

The intuition for the first part of Proposition 3 is trivial: if cost is high enough, then store b's optimal strategy is to exit. The less obvious part of the result is that, for a given c, exit takes place for a high enough s. The idea is that, as the discussion of Proposition 1 makes clear, an increase

Fig. 3. Optimal choice in the (s, c) space (number of titles, fixed cost) in the linear case
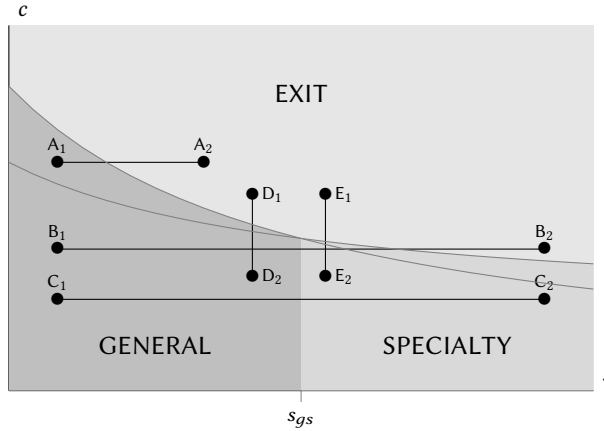
in s makes store a relatively more attractive, and thus reduces store b's profit. The condition in Proposition 3 is required because, if c is low enough, then a specialty store makes a positive profit regardless of the value of s. In other words, a specialty store's profit converges to a positive lower bound as s tends to infinity.

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 s and c. Moves A, B and C correspond to an increase in the number of titles, s. 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 s increases, we observe a general store exit. (Recall that, if s 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 s is low. As s increases, long after store A has gone out of business, B remains active, but past $s = s_{gs}$ becomes a specialty store. As s 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 s is, it remains active. Notice however that, similarly to B, store C becomes a specialty store when $s > s_{gs}$.

Moves D and E correspond to a decrease in c. In case D, we observe the entry of a general store, whereas in case E we observe the entry of a specialty store. Naturally, the move would be reversed if we considered an increase in c. 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). Consider

for example a decrease in c (more favorable local demand conditions). At a time when s is low, the new entrant would have entered as a general store. However, as s increases, the same decrease in c is now more likely to lead to the entry of a specialty store.

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 s is close to the disruption level s_{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 s 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 s 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.

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 s 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). 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 a corner solution where a specialty store is always optimal).

Assumption 2. [*No corner solution*] $p_a > \overline{z} + \frac{m(k) - \sqrt{2}\, m(k/2)}{\sqrt{2} - 1}$.

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. [*Specialty-store with endogenous pricing*] *Suppose Assumption 2 holds. There exists a threshold s_{gs} such that store b optimally chooses to be a specialty store if $s > s_{gs}$. In the right*

neighborhood of s_{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 s), then firm b is better off by becoming a specialty store. The main intuition for the s-threshold part of Proposition 4 is similar to Proposition 1: As total supply s 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 s_{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.

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 $m(s/2)$. By contrast, an eclectic consumer values the online store at $m(s)$ 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 $m(s/2)$. By contrast, an eclectic consumer values the online store at $m(s)$ whereas the general store is 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, $m(s/2) - m(k/2)$ or $m(s) - m(k)$. Notice that $m(s) - m(k) > m(s/2) - m(k/2)$ if and only if $m(s) - m(s/2) > m(k) - m(k/2)$. Since $s > k$, $s - s/2 > k - k/2$, which would suggest the inequality holds. However, concavity of $m(x)$ would work against the inequality. Suppose that $F = v$ is linear, so that $m(x) = x/(1 + x)$. Then the function $m(x) - m(x/2)$ is non-monotonic, first increasing for $x \in [0, \sqrt{2}]$ and then decreasing. This implies that we can find values of s 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 the latter, due their limited

size, cannot 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.

■ **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, stores b_0 and b_1, do compete for the same potential demand. Specifically, we assume a consumer is characterized by a value z and a relative preference between stores b_0 and b_1 in the form of a location $d \in [0, 1]$ and transportation cost t per unit of distance to store b_0 (located at 0) and to store b_1 (located at 1). Moreover, we assume that d and z are independently and uniformly distributed: $d \sim U[0, 1]$ and $z \sim U[0, \overline{z}]$. Our main result is that, under competition, the genre choice exhibits strategic complementarities.

Proposition 5. *[Strategic complementarity in specialization] Let s be such that store b_0 and b_1 are indifferent between being general store and being a specialty store. In the neighborhood of s, being a specialty store is a strict best response to the rival choosing to be a specialty store.*

Proposition 5 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 s crosses a certain threshold, say $s_\circ$, 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 s is lower than $s_\circ$* (by a little), then firm b_1 also optimally switches to being a specialist (specializing in the niche that firm b_1 did not).

In sum, Proposition 5 provides an additional force in the direction of specialization.

To conclude this section, we note how Amazon is strictly worse off when competing with two specialty stores compared to two generalist stores. Again, this suggests caution when interpreting a higher degree of consumer polarization as a desirable outcome for larger, online retailers.

■ **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 $m(s/2) - m(k/2)$. It follows that only consumers with a value of z greater than z'' purchase at the bricks-and-mortar store. Since z is uniformly distributed, we conclude that firm b's market share is given by $q_g = \overline{z} - z''$.

Consider now the case of a specialty store b. Its disadvantage with respect to store a is given by $m(s/2) - m(k)$. It follows that only consumers with a value of z greater than z' purchase at

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

the bricks-and-mortar store. Since z is uniformly distributed, we conclude that firm b's market share (among its genre followers) is given by $q_s = \overline{z} - z'$. 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 $(\overline{z} - z')/2$.

The values of s and k were selected so that $\pi_g = \overline{z} - z'' = (\overline{z} - z')/2 = \pi_s$. In other words, for the particular values of s 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

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

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 6. *[Welfare] 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

Our paper is primarily based on a theoretical exercise. However, we believe it has 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, thus avoiding broad genres on which Amazon's advantage is hard to counteract. Second, Daunt stresses how increasing offline perks is another key to differentiation. Indeed, our model shows that small increases in these perks can be as profitable as large increases in store assortment.

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, more and more brick and mortar stores will be forced to exit, as the generalist strategy (the only one effectively available to them in this scenario) becomes unprofitable.

■ **The tyranny of 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.

■ **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.** 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 increasing 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.

■ **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 and 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 totally fail to replicate even a small fraction of Amazon's variety. Since our analysis provides a rationale for the growth in the number of independent bookstores (in the US and in recent years), it also provides support for Bookshop's strategy.

■ **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 2021. 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 the next Section 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. We also have detailed information on the approximately 2,800 bookstores, including type of store and address, which

Table 1. Aggregate data by channel

	Chains	Bookstores	Mass Mer.	Online D2C
(a) # titles	43,887	39,267	12,875	47,903
(b) # books	127,602,337	31,701,747	171,420,650	163,995,077
(b)/(a)	2,907	807	13,314	3,423

we have matched to publicly available geographic and demographic data. Specifically, we divide bookstore orders into four different channels:

- **Online D2C**: Sales made to Amazon.
- **Bookstores**: Sales made to independent bookstores (an aggregated version of the bookstore level data).
- **Bookchains**: Sales made to bookstore chains such as Barnes & Noble etc.
- **Mass Merchandiser**: Sales made through large non-specialty stores such as Target, Walmart etc.

Since purchases are rather sparse (i.e., there are many zeros), and since individual bookstores rarely reorder the same book over multiple months, we aggregate orders at the title-author level, over time (2016-2019), and across multiple stores owned by the same firm. This results in a sample of 39,000 unique book titles purchased, for a total of over 5,700,000 transactions.[9]

We now present a variety of facts that corroborate our theoretical findings.

■ **Stocking decisions across channels.** 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 bricks-and-mortar stores' relatively small size. In other words, Proposition 2 uncovers a novel reason why Amazon's growth indirectly leads to a thickening of the long tail.

One simple way to test these predictions is to compute concentration indexes by type of channel. To this end, we first compute the number of books and titles ordered by different channels. Then, we ask: What does the distribution of sales look like? How does it differ across channels? To answer this question we compute the percentage of sales due to the top N books.

Table 1 shows that, despite being by far the smallest channel in terms of book orders, bookstores combine for nearly as many title orders as chains and Amazon, and over three times as many title orders as mass merchandisers. This offers initial, suggestive evidence of bookstores' shying away from a generalist strategy. If each bookstore was a generalist, they would also be quite homogeneous. But then, given their limited size (the average bookstore in our dataset orders around 1000 titles), the total number of titles purchased by US bookstores would be nowhere close to 39267. At the same time, the average number of books sold per title would be considerably higher.

■ **The offline long tail.** We now turn to studying the sales distribution across different channels. Table 2 shows, for multiple values of N, the percentage of sales accounted for by the (channel specific) top N sellers. Consistent with Proposition 2, the percentage of sales corresponding to the top N titles is lower at bookstores — both chain stores and independent ones — than it is at Amazon. This remains true even for large values of N. For instance, $N = 10,000$ is about 10 times

---

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

Table 2. Sales concentration by channel

N	Book Chains	Book Stores	Mass Merchand.	Online D2C
100	11.2	11.1	21.4	14.7
500	28.7	26.0	54.4	34.0
1000	39.9	36.1	71.2	45.7
2500	58.0	53.1	89.4	62.8
5000	72.5	68.2	97.7	75.8
7500	80.9	77.2	99.4	82.9
10000	86.6	83.3	99.9	87.4

Table 3. Niche genres

Bookstores	4.6
Chains	2.7
Online	3.7
Mass Merchants	0.6

the size of an average bookstore; nevertheless, bookstores' specialization on (a variety of) niches limits the percentage of sales the top 10,000 books account for.

Table 2 and Proposition 2 challenge the Anderson [2004] view that the long tail is an online phenomenon, that is, the prediction that "the Internet channel exhibits a significantly less concentrated sales distribution when compared with traditional channels" [Brynjolfsson, Hu, and Simester, 2011, p. 1373].

Finally, it is interesting to see how these figures are dramatically higher for mass merchandisers: by their very definitions, these stores tend to be quite homogeneous across the US, and concentrate their sales on a relatively limited set of popular books (the top 1000 sellers on this channel account for around 71% of its total sales, around twice the equivalent figure for bookstores). So, while the Anderson [2004] intuition captures the Amazon vs mass merchandisers dichotomy quite well, we find that it falls short of explaining the low concentration of sales displayed by other offline retailers.

■ **Niche genres.** Much of our analysis refers to niche genres. We now dig deeper into this issue. We define niche genres as below-median genre market share. All together, niche genres so defined account for a *combined* market share slightly lower than 2.7%. To corroborate our theory that bricks-and-mortar stores specialize in narrow niches as a result of Amazon's growth, we look at niche sales by channel.[10]

Table 3 shows the percentage of sales accounted for by nice genres, classified by type of retailer. The values confirm the idea that bricks-and-mortar bookstores sell the highest percentage of niche genres: around 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).

---

10. A related note: While niche genres and niche titles are distinct categories, they are correlated: Titles in the bottom quintile of sales are 9% more likely to be of a niche genre.

Table 4. Share of niche sales

Urban vs non-urban	+47%
Small urban vs large urban	+46%
Small non-urban vs large non-urban	+116%

Our model offers an additional prediction: small bookstores have stronger incentives to follow a niche strategy compared to larger ones. In our data, however, one potential confounding factor arises: small bookstores are more likely to have urban locations. An alternative interpretations for the number in Table 3 would be that urban consumers have, on average, a stronger taste for niche books. The argument might be, for example, that urban consumers are on average more educated and thus more likely to have formed an interest in specialized subjects such as astronomy, machine learning, or European art history.

In order to address this alternative explanation, we present our results for urban and rural bookstores separately. Moreover, we split the stores into small and large stores, so as to explicitly consider our prediction that a niche strategy is more likely to be followed by small stores. Specifically, we define small bookstores all of those who order fewer than 300 books (the median is around 1700 books).

Table 4 shows the results of this alternative tabulation. The first row confirms our intuition that an urban-rural divide is present. However, as the following two rows show, even controlling for this gap, it is still overwhelmingly the case that small bookstores are more likely to specialize on niche genres, as predicted by our theory.

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 Andersons'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 impact on the rise of the offline long tail.

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.

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," Harvard Business School Working Paper 20-068.

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, illustrated edition edition.

Appendix

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

$$z + m(k/2)$$

By contrast, buying at a yields expected value

$$m(s/2)$$

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

$$z = m(s/2) - m(k/2)$$

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

$$\pi_g = 1 - \left(m(s/2) - m(k/2) \right) / \overline{z} \tag{1}$$

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

$$z + m(k)$$

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

$$m(s/2)$$

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

$$z = m(s/2) - m(k)$$

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

$$\pi_s = \tfrac{1}{2} \left(1 - \left(m(s/2) - m(k) \right) / \overline{z} \right) \tag{2}$$

(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 $s = 0$, 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 s_g is such that $\left(m(s_g/2) - m(k/2) \right) / \overline{z} = 1$. For $s = s_g$, we have $\pi_g = 0$, whereas

$$\pi_s = \tfrac{1}{2} \left(1 - \left(m(s_g/2) - m(k) \right) / \overline{z} \right) > \tfrac{1}{2} \left(1 - \left(m(s_g/2) - m(k/2) \right) / \overline{z} \right) = 0$$

Such an s will exist whenever $\lim_{s \to \infty} \left(m(s/2) - m(k/2) \right) / \overline{z} > 1$, which is implied by Assumption 1. (As mentioned in the text, if this condition does not hold — for instance because $\overline{z}$ 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 $s_{gs} \in (0, s_g)$ such that $\pi_g(s_{gs}) = \pi_s(s_{gs})$, where for notational simplicity we have suppressed the store profit's dependence on k and $\overline{z}$. To show that s_{gs} is unique we note that

$$\frac{d \left(\pi_s - \pi_g \right)}{ds} = \left(-m'(s) + 2 m'(s) \right) / (4 \overline{z}) = m'(s)/(4 \overline{z}) > 0 \tag{3}$$

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

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

$$\frac{\partial(\pi_s - \pi_g)}{\partial k} = \frac{m'(k)}{2\,\overline{z}} - \frac{m'(k/2)}{2\,\overline{z}} = \frac{1}{2\,\overline{z}}\left(m'(k) - m'(k/2)\right) < 0 \tag{4}$$

where the inequality follows from concavity of m [David, 1997]. Similarly,

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

where the second equality follows from (1) and (2). By definition, $\pi_s = \pi_g = \overline{\pi}$ at $s = s_{gs}$. It follows that

$$\left.\frac{\partial(\pi_s - \pi_g)}{\partial\overline{z}}\right|_{s = s_{gs}} = (\tfrac{1}{2} - \overline{\pi})/\overline{z} - (1 - \overline{\pi})/\overline{z} = -1/(2\,\overline{z}) < 0 \tag{5}$$

By the implicit function theorem,

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

where the inequality follows from (3) and (4). Also by the implicit function theorem,

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

where the inequality follows from (3) and (5). ∎

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

$$m(\alpha s) = m(k)$$

where αs 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 αk is interested in the genre offered by store b, and a fraction $(m(\alpha s) - m(k))/\overline{z}$ of this fraction prefers store b to store a. This implies that store b's profit from specializing in genre x is given by

$$\pi_x = \alpha k\left(1 - (m(\alpha s) - m(k))/\overline{z}\right)$$

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

$$\pi_y = (1 - \alpha)\left(1 - \left(m((1 - \alpha)s) - m(k)\right)/\overline{z}\right)$$

If $s = 0$, 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 s_x be the value of s such that $\pi_x = 0$. Such an s will exist whenever $\lim_{s\to\infty}(m(\alpha s) - m(k))/\overline{z} > 1$, which is equivalent to Assumption 1. We then have

$$\pi_y = (1 - \alpha)\left(1 - \left(m((1 - \alpha)s_x) - m(k)\right)/\overline{z}\right) > \alpha k\left(1 - (m(\alpha s_x) - m(k))/\overline{z}\right) = 0$$

(If this condition does not hold — for instance because $\overline{z}$ 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 $\widehat{s}_{xy} \in (0, s_x)$ such that $\pi_g(\widehat{s}_{xy}) = \pi_s(\widehat{s}_{xy})$, where for notational simplicity we have suppressed the store profit's dependence on k and $\overline{z}$. Let s_{xy} be the highest of these values. Then $\pi_y \geq \pi_x$ for $s > s_{xy}$. ■

Proof of Proposition 3: A general store net profit is given by

$$\pi_g = k \left(1 - (m(s/2) - m(k/2)) / \overline{z} \right) - c\,k$$

It follows exiting is better than being a general store if and only if

$$c > c_g(s) \equiv 1 - (m(s/2) - m(k/2)) / \overline{z}$$

Note that

$$\frac{d\,c_g(s)}{d\,s} = -m'(s/2)/\overline{z} < 0$$

A specialty store's net profit is given by

$$\pi_s = \tfrac{1}{2} k \left(1 - (m(s/2) - m(k)) / \overline{z} \right) - c\,k$$

It follows that exiting is better than being a specialty store if and only if

$$c > c_n(s) \equiv \tfrac{1}{2} \left(1 - (m(s/2) - m(k)) / \overline{z} \right)$$

Note that

$$\frac{d\,c_n(s)}{d\,s} = -\tfrac{1}{2} m'(s/2)/\overline{z} < 0$$

Proposition 1 implies that $c_g(s) > c_n(s)$ if and only if $s < s_{gs}$, where s_{gs} is the critical value (derived in Proposition 1) such that $\pi_g = \pi_s$. (Figure 3 illustrates the result in the case when $F(v) = v$.) It follows that

$$c'(s) \equiv \min \left\{ c_g(s), c_n(s) \right\}$$

defines a downward-sloping boundary such that exit is optimal if and only if $c > c'(s)$. Taking limits, we find that

$$\lim_{s \to \infty} \pi_s(s) = \tfrac{1}{2} k \left(1 - (\overline{v} - m(2\,k)) / \overline{z} \right) - c\,k$$

This is positive if and only if $c < \overline{c}$. It follows that if $c < \overline{c}$, then exit never takes place, whereas if $c > \overline{c}$ there exists a finite $s^\circ(c)$ such that exit takes place if and only if $s > s^\circ(c)$. ■

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

$$q_g = 1 - \left(m(s/2) - m(k/2) - p_a + p_g \right) / \overline{z}$$

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

$$\widehat{p}_g = \tfrac{1}{2} \left(\overline{z} - m(s/2) + m(k/2) + p_a \right) \tag{6}$$
$$\widehat{q}_g = \tfrac{1}{2} \left(\overline{z} - m(s/2) + m(k/2) + p_a \right) / \overline{z} = \widehat{p}_g / \overline{z} \tag{7}$$
$$\widehat{\pi}_g = \widehat{p}_g\,\widehat{q}_g = (\widehat{p}_g)^2 / \overline{z} \tag{8}$$

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

$$q_s = \tfrac{1}{2} \left(1 - (m(s/2) - m(k) - p_a + p_s) / \overline{z} \right)$$

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

$$\widehat{p}_s = \tfrac{1}{2}\left(\overline{z} - m(s/2) + m(k) + p_a\right) \tag{9}$$

$$\widehat{q}_s = \tfrac{1}{4}\left(\overline{z} - m(s/2) + m(k) + p_a\right)/\overline{z} = \widehat{p}_s/(2\,\overline{z}) \tag{10}$$

$$\widehat{\pi}_s = \widehat{p}_s\,\widehat{q}_s = (\widehat{p}_s)^2/(2\,\overline{z}) \tag{11}$$

Direct inspection of (6) and (9) reveals that

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

that is, in equilibrium specialty bookstores set a higher price. Moreover, from (6)–(7) and (9)–(10) we conclude that

$$\widehat{p}_s/\widehat{q}_s = 2\,\overline{z} > \widehat{p}_g/\widehat{q}_g = \overline{z} \tag{12}$$

Consider the extreme case when $s = 0$. Straightforward computation shows that $\widehat{\pi}_g > \widehat{\pi}_s$ if and only if Assumption 2 holds. At the opposite end, let s_g be such that $\widehat{p}_g = 0$. Comparing (6) and (9), we see that, at $s = s_g$, $\widehat{p}_s > \widehat{p}_g = 0$. From (8) and (11) we conclude that, at $s = s_g$, $\widehat{\pi}_s > \widehat{\pi}_g = 0$. Since both $\widehat{\pi}_s$ and $\widehat{\pi}_g$ are continuous we conclude by the intermediate-value theorem that there exists at least one $\widetilde{s}_{gs}$ such that $\widehat{\pi}_s = \widehat{\pi}_g$. Let s_{gs} be the highest of these values. Then $\widehat{\pi}_s > \widehat{\pi}_g$ when $s_{gs} < s < s_g$.

Finally, notice that, at $s = s_{gs}$, $\widehat{\pi}_g = \widehat{\pi}_s$, that is, $\widehat{p}_g\,\widehat{q}_g = \widehat{p}_s\,\widehat{q}_s$. Since, from (12), $\widehat{p}_s/\widehat{q}_s > \widehat{p}_g/\widehat{q}_g$, it must be that, at $s = s_{gs}$, $\widehat{p}_s > \widehat{p}_g$ and $\widehat{q}_s < \widehat{q}_g$. Since these are strict inequalities, they also hold in the neighborhood of $s = s_{gs}$. It follows that, in the right neighborhood of $s = s_{gs}$, a specialty store earns a higher profit, sets a higher price, and captures a lower market share. ∎

Proof of Proposition 5: Figure 5 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: $d \sim U[0, 1]$ and $z \sim U[0, \overline{z}]$. 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 5 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

$$m(k/2) + z - t\,d > m(s/2)$$

whereas beating store b_1 requires

$$m(k/2) + z - t\,d > m(k/2) + z - t\,(1 - d)$$

This results in the following set of inequalities

$$z > m(s/2) - m(k/2) + t\,d$$
$$d < \tfrac{1}{2}$$

which in turn correspond to the areas in blue (top panel) and red (bottom panel).
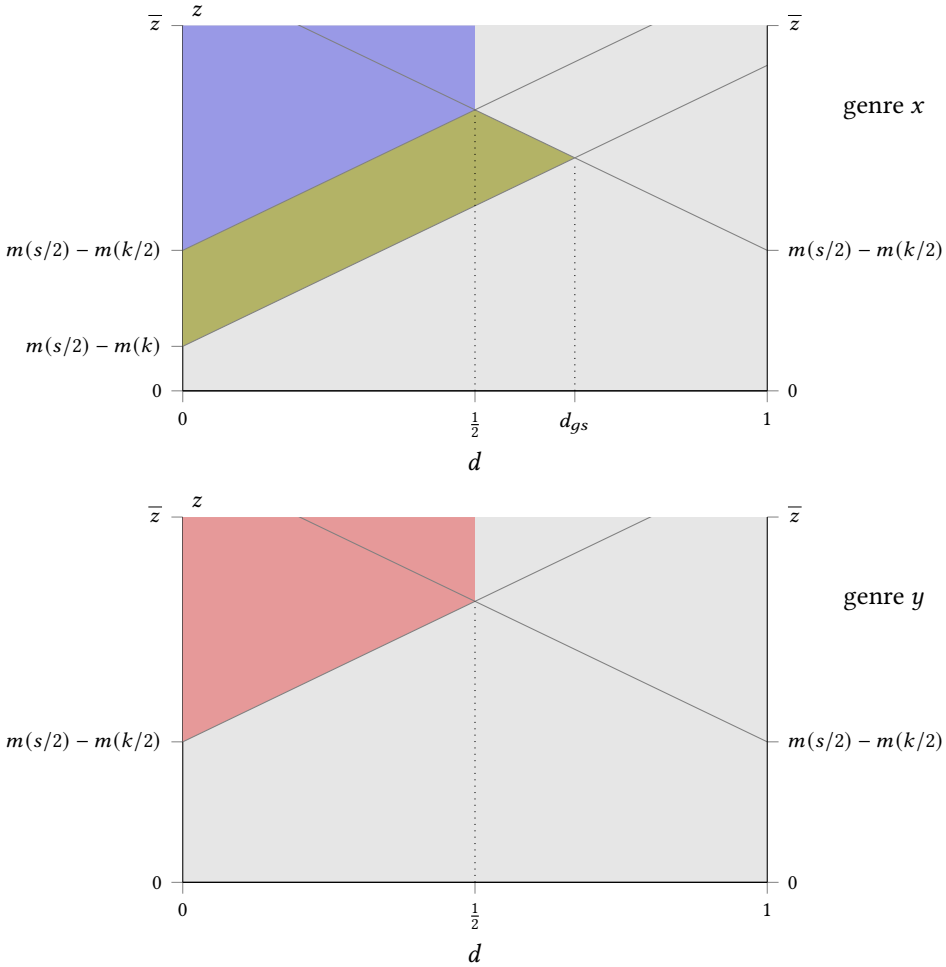
Given that b_1 chooses to be a general store, how does b_0 change its profits by specializing in genre x? Store b_1's demand from x consumers is now determined by

$$m(k) + z - t\,d > m(s/2)$$

(beat firm a) and

$$m(k) + z - t\,d > m(k/2) + z - t\,(1 - d)$$

Fig. 5. Store strategy under bricks-and-mortar competition

(beat firm b_1). This simplifies to

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

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 s and k in Figure 5 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 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
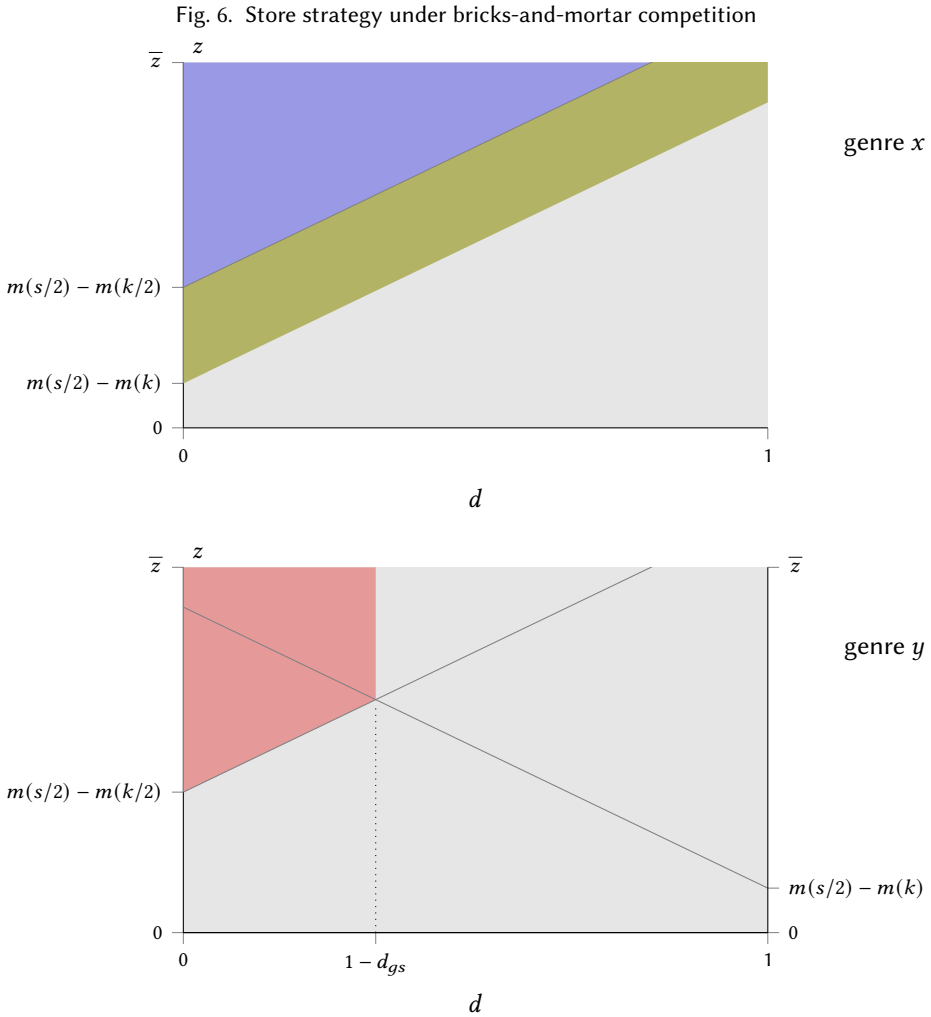
Fig. 6. Store strategy under bricks-and-mortar competition

Figure 6. 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

$$m(k/2) + z - t\,d > m(s/2)$$

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

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

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 6 is greater than the green area in the top panel of Figure 5, whereas the red area in the bottom panel of Figure 6 is lower than the red area in the bottom panel of Figure 5. 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. ■