Quantitative Marketing and Economics (2020) 18: 125–154
https://doi.org/10.1007/s11129-019-09214-6

Retailers' product location problem with consumer search

Raluca M. Ursu[1] · **Daria Dzyabura**[2]

Received: 15 November 2018 / Accepted: 26 June 2019 / Published online: 12 September 2019
© Springer Science+Business Media, LLC, part of Springer Nature 2019

Abstract

With few exceptions, today's retailers sell products across multiple categories. One strategic consideration of such retailers is product location, which determines how easy or difficult different categories are for customers to access. For example, grocery or department stores determine which products will be located closer to the entrance of the store versus at the back of it, while online retailers decide which products to feature on the homepage, and which will require scrolling or keyword search to get to. In this paper, we study how a retailer should optimally locate products within a store, when the locations chosen affect consumer search costs. We show that the retailer has an incentive to prioritize products with lower utility, contrasting with prior work. The intuition for our result is that the consumer may be willing to search less preferred products only at the lower cost, while the more preferred products will be searched even at higher search costs. This strategy benefits the retailer by increasing the number of products the consumer searches and thus, the ones she may buy. Our finding is robust to several extensions: (i) a retailer determining not only product locations, but also prices, (ii) independent (e.g. categories), as well as substitute products, and (iii) a focal retailer that faces competition. From a managerial perspective, we show that allocating products in the store without taking into account how this affects consumer search costs, might mean consumers overlook products they would otherwise purchase.

Keywords Consumer search · Multi-category retailer · Product location problem

JEL Classification L81 · D83 · D11

✉ Raluca M. Ursu
rursu@stern.nyu.edu

Daria Dzyabura
ddzyabura@nes.ru

[1] Marketing Department, Stern School of Business, New York University, 40 West 4th Street, New York, NY 10012, USA

[2] New Economic School, Skolkovskoe Shosse, 45, Moscow 143025, Russia

1 Introduction

With few exceptions, today's retailers sell products across multiple categories. Examples include grocery stores selling products in categories such as dairy, frozen foods, produce and flowers, or online retailers like Amazon.com selling products across more than 30 categories. One strategic consideration of such retailers is product location, which determines how easy or difficult different categories are for customers to access. For example, grocery stores must allocate different product categories to different aisles within a store, such as closer to the entrance of the store or near the back of it. Even within a single aisle, some products can be placed at eye level, while others are placed on low shelves near the floor, making the latter more costly for consumers to search. Yet other categories are placed near the checkout, where the cost of searching them is very low. Similarly, an online retailer can choose some categories to feature on the homepage, making them relatively easier (less costly) to access than categories that require scrolling through several menus, or using keyword search to get to.

In this paper, we study how a retailer should optimally allocate products within a store, when the locations chosen affect consumer search costs. To this end, we develop a model in which a retailer sells independent products (e.g. categories), that are neither complements nor substitutes. Examples include product categories such as flowers or milk in a grocery store, and perfume or dress shirts in a department store. The retailer has access to a set of locations within a store, which have different, exogenously defined, search costs. The retailer chooses how to display products within a store to maximize its expected payoff, thereby determining a product's search cost. Taking product location and thus search costs as given, a consumer enters the store with a belief about the potential (expected) utility she can obtain from each product. For example, she knows whether she values milk or flowers more on a certain grocery shopping trip, regardless of the product's location. By paying a cost, the consumer searches a product, observes its realized utility, and decides whether or not to purchase it. The goal of the consumer is to maximize her total utility net of search costs by optimally determining which subset of products to search and whether or not to purchase a searched product.

In this setting, we derive the optimal solution for the retailer's product location problem. First, we show analytically for the case of two products, that the retailer has an incentive to prioritize the product with lower expected utility. That is, to maximize its expected payoff, the retailer will place the lower expected utility product in the lower search cost location, subject to the constraint that both products' expected utility is high enough that the consumer will search them in these locations. The intuition for our result is that the consumer may be willing to search less preferred products only at the lower cost, while the more preferred products will be searched even at higher search costs. Thus, the proposed strategy may increase the number of products the consumer searches and then buys, thereby increasing the expected payoff of the retailer. The stylized fact that most grocery retailers place milk (a high utility category) in the back of the store, while placing less necessary items by the entrance and check out provides validity to this result.[1]

Second, we generalize our solution by constructing the optimal allocation algorithm for any number of products. Again, we see that the algorithm prioritizes lower expected utility products, placing them in locations with lower search costs.

Our findings contribute to and contrast with those in the literature on ranking algorithms for information retrieval. This literature focuses on the problem faced by a search engine, which tries to locate the most relevant search result (e.g. a website URL) at the top of a ranked list. One reason for the contrast is that, while search engines compete solely on the quality of their ranked lists, retailers may compete on several other dimensions, such as physical location or product assortment, allowing them to be strategic about product location. Thus, a retailer can choose the location of products within the store strategically in order to maximize profits.

In addition to the main analysis, we extend our baseline model in several directions to show the robustness of our result. First, we show that our result continues to hold when the retailer chooses optimally not only how to allocate products, but also each product's price. In this model, the consumer searches both for a product's utility and for its price. The retailer faces a tradeoff: on the one hand, placing the high expected utility product in a low cost location means being able to charge a higher price for it (both because of its higher expected utility and because of its lower search cost), but risking that the consumer will not search the low expected utility product; on the other hand, placing the low expected utility product in a low cost location might involve charging a lower price for it, but ensuring that the high expected utility product will be searched. Once again, we show that the retailer has an incentive to prioritize the lower expected utility product in the low cost location.

Second, we consider the case where products are substitutes, rather than independent, that is the consumer will purchase at most one of the two products (i.e. unit demand) on a given shopping trip. This setting is the most common one found in the theoretical search literature (e.g. Weitzman 1979), modeling the case where the consumer is searching for the best alternative within a product category. In this case, unlike when products are independent, what the consumer observes while searching one product affects her decision to search another product, to stop searching, and to make a purchase decision. We again find that the retailer has an incentive to prioritize the lower expected utility product in the low search cost location. This is the case because by prioritizing the lower expected utility product, the consumer is more likely to search the other product (if the first search is unsatisfactory), and thus to purchase one of them, compared to the case where the higher expected utility product is prioritized and the consumer might stop after just one product searched. Thus,

[1] Milk, as well as other products, when on sale, may serve as loss-leader products, encouraging consumers to enter the store (Johnson 2017). Distinct from this mechanism, we provide a novel rationale for placing high expected utility items in the back of the store that is unrelated to their price and does not require them to be sold below cost.

we show that prioritizing lower expected utility options can benefit the retailer in the case of product substitutes as well.

Finally, we study the impact of competition on the optimal product location problem of the focal retailer. We show that when retailers compete solely in their product allocation, then prioritizing higher expected utility products is optimal. However, in reality retailers may be differentiated in several dimensions, therefore not competing directly in their product allocation. For example, retailers may differ in terms of their physical location (requiring transportation costs to visit multiple retailers), their product assortments, or other dimensions, such as customer service or loyalty programs, which increase consumer switching costs between retailers. When this is the case, we show that the optimal product location solution is still to place the lower expected utility product in the low search cost location.

Our results emphasize the need to match a product's expected utility to its search cost in order to maximize retailer payoffs. As such, our results have managerial implications for both offline grocery and department stores and online retailers choosing how to optimally allocate products to locations that affect the ease with which consumers search these products. Allocating products without taking into account how their location affects consumer search costs, might mean consumers overlook products they would otherwise purchase.

The remainder of this paper is organized as follows. The next section reviews relevant prior work. Section 3 introduces our model, while Section 4 presents our main results. In Section 5 we describe extensions of the model, while the last section concludes.

2 Related literature

This paper relates primarily to three strands of the literature: (i) the theoretical consumer search literature, (ii) the literature on in-store product location strategy, and (iii) prior work on designing optimal ranking algorithms for information retrieval. In what follows, we describe how our paper relates and contributes to the previous literature.

First, this paper is related to theoretical work on consumer search. Prior work focuses on characterizing optimal consumer search decisions in different settings, for example when determining which products to search (Stigler 1961; Weitzman 1979), which product attributes to obtain information about (Branco et al. 2012, 2016), or which products to search more intensively than others (Chick and Frazier 2012; Ke et al. 2016; Ke and Villas-Boas 2017). We contribute to this literature by studying a new setting in which consumers search optimally across multiple product categories. Closely related is also work on ordered search where consumers inspect options in a predetermined order (Arbatskaya 2007; Wilson 2010; Zhou 2011; Rhodes 2011; Armstrong et al. 2009; Armstrong and Zhou 2011; Haan et al. 2018; Gamp 2017; Petrikaite 2018). For instance, the work of Gamp (2017) and Petrikaite (2018) relates to ours in that the authors consider a model where a firm affects product search order by choosing the magnitude of consumer search costs for a product. In contrast, in our paper search costs are exogenous, that is they are determined by the layout of the store, and the firm chooses which products to allocate to which location. Moreover, in terms of findings, our paper is also related to models of intermediary search diversion. Hagiu and Jullien (2011) present a model where an intermediary, who has more information than consumers about product matches, may benefit from diverting a consumer away from her preferred option, if this strategy influences the type of consumers visiting the store or sellers' pricing decisions. We contribute to this literature by showing that a similar result arises in a setting where the retailer, with less information about the consumer, chooses how to allocate products within the store, taking into account that these locations may affect consumer search costs. Finally, our paper is related to theoretical work on multi-category search (Burdett and Malueg 1981; Carlson and McAfee 1984; Anglin 1990; McAfee 1995; Gatti 1999; Shelegia 2011; Zhou 2014). This research models how a consumer, searching for several products, decides which subset of them to buy from one of many retailers she visits. In contrast, we consider the problem of a consumer searching for multiple products at a single retailer.

Second, our paper relates to work on firm in-store product location strategy. One stream of this work shows that the location of products in the store affects consumer attention and purchase decisions. For example, Chandon et al. (2009) show that placing products at the top of a shelf increases attention and choice. Larson et al. (2005) identified the most common shopping paths in a supermarket and showed that consumers spend most time traveling on the outer ring of the store, making products placed at the ends rather than the center of an aisle more salient. Similarly, products near the checkout are more salient since almost all consumers pass by that area. These findings lend support to our core assumption that different locations in the store have different search costs, affecting how consumers search and what products they purchase. We contribute to this stream of research by considering the optimal search problem of the consumer in a store environment and how this affects the retailer's product allocation strategy. Work on price obfuscation and retail strategies to increase search frictions is also related to ours (Ellison and Ellison 2009; Ngwe et al. 2019). Such work focuses on how obfuscation allows retailers to price discriminate heterogeneous consumers, by discouraging less price sensitive ones from further search, leading them to buy more expensive products. In contrast, in our paper the retailer can benefit from prioritizing lower utility options even in the absence of consumer heterogeneity. In addition, our paper is also related to research studying the relation between in-store travel distance and unplanned spending. Previous work has shown that increasing in-store travel distance, either through relocating products (Granbois 1968) or through strategic product promotions (Hui et al. 2013), can increase unplanned spending. The intuition for this result is similar to that for our findings: by encouraging consumers to be exposed to more product categories (by traveling longer distances), they may purchase from more of them. Our contribution to this literature is the theoretical treatment of the problem of how to allocate products within the store given optimal consumer search.

Third, by interpreting ranked lists as ordered locations in which products can be placed, our paper is also related to work on optimal algorithms for ranking a set of options online. Research in marketing and economics studies the impact of ranked lists on consumer choices (Ghose et al. 2012a, b; Ghose et al. 2014; Chen and Yao 2016; De los Santos and Koulayev 2017; Ursu 2018). For example, Ursu (2018) shows that even when the ranked list is randomly generated, the position of a product affects consumer choices. In addition, work in computer science and marketing studies a search engine's problem of how to construct product rankings for information retrieval (e.g. Liu 2011). Algorithms, such as 'learning to rank', score the likely relevance of a search result for a consumer (e.g. the likelihood of a purchase or a search, measured using metrics such as normalized discounted cumulative gain) and then rank the option with the highest score first (e.g. Yoganarasimhan 2018). Such algorithms are optimal for search engines that compete solely in the quality of their ranked lists. In contrast, our paper studies the problem of a retailer, that due to higher product differentiation, on factors such as physical location, product assortments, customer service, or loyalty programs, may compete on several dimensions other than just the allocation of products within the store. Also, location mainly affects a product's accessibility, rather than consumer expected utilities. As a result, we show that this retailer can be strategic about product location and may choose not to prioritize the highest utility product for a consumer, as in the case of grocery stores choosing to place milk (a high utility product) at the back of the store.

In addition, our work is related to papers considering the design of sponsored ad auctions in the presence of consumer search (Varian 2007; Chen and He 2011; Athey and Ellison 2011). The result of the auction is an ordered list of options consumers can choose from, and is thus related to our work if we interpret this list as a set of locations where products can be placed. We differ from this work in several ways, most importantly in studying the product allocation problem of a retailer, rather than the bidding strategies of individual advertisers, and in allowing different locations to have different search costs.

Finally, by studying a multi-category retailer, our paper is related to previous work on how consumers make multi-category purchases (Ainslie and Rossi 1998; Chintagunta and Halder 1998; Manchanda et al. 1999; Erdem and Winer 1999; Seetharaman et al. 1999, 2005; Chib et al. 2002; Singh et al. 2005; Hansen et al. 2006; Song and Chintagunta 2006, 2007; Mehta 2007). We contribute to this literature by studying not only multi-category purchases, but also multi-category search, as well as optimal retailer decisions in such settings.

3 Model setup

In this section, we present our baseline model. Then, in Section 5 we show that our main results are generally robust to several extensions.

A retailer has access to two products $j \in J = \{1, 2\}$, that are independent (e.g. categories), that is neither substitutes nor complements. It chooses a location for each product in the store, indexed by $l \in L = \{1, 2\}$. Let j_l denote the product in location l. These locations have different, exogenously defined, search costs. A representative consumer enters the store with a belief distribution about the potential (expected) utility she can obtain from each product. For example, she knows whether she values milk or flowers more on a certain grocery shopping trip, regardless of the product's location. By paying a cost, the consumer searches a product, observes its realized utility, and decides whether or not to purchase it. The consumer may purchase several products in this setup. In Section 5, we also consider the case of product substitutes, where the consumer purchases at most one product, and the case where there are more than two products available.

Consumer search problem A consumer searches the products optimally. Searching a product j reveals all uncertainty about it, giving the consumer a potential utility ε_j assumed to be drawn from a product specific distribution function $F_j(\cdot)$ and defined on the interval $[\underline{\theta}_j, \overline{\theta}_j]$.[2] This utility is the consumer's match value with a product and it cannot be observed by the retailer. However, both the consumer and the retailer know the distribution of match values for all products, and thus have a belief about the expected utility from a product. Because products are independent, the consumer has product specific outside options, representing her expected utility from not buying each product, which we normalize to zero.

Searching a product is costly, and this cost is made up of two components. The first component, c_{j0}, is the baseline cost of searching the product once the consumer has arrived at j's location. The second component depends on the product's location in the store. Some locations are costlier to search than others. We number the locations in increasing order of search cost, such that the first location is the least costly to search, and the last location is the most costly to search. For simplicity, we assume the cost differences between subsequent locations are constant, that is, the location cost takes value $\alpha(l-1)$ for product location l. The net cost of searching the product j_l in location l is given by

$$c_{j_l} = c_{j0} + \alpha(l-1). \tag{1}$$

The locations of the products are chosen by the retailer and cannot be changed by the consumer. Given the locations of the products, the consumer decides optimally which products to search by computing the expected utility from searching each product. She only searches products whose expected utility of searching exceeds the cost of searching, or more formally those for which

$$\int_0^{\overline{\theta}_{j_l}} \varepsilon dF_{j_l}(\varepsilon) - c_{j_l} \geq 0. \tag{2}$$

Conditional on search, the consumer purchases j_l if its realized utility is greater than the outside option, that is, she buys if $\varepsilon_{j_l} \geq 0$. Denote this purchase probability by $\phi_{j_l} = 1 - F_{j_l}(0)$.[3]

Let $\gamma_j = \int_0^{\overline{\theta}_j} \varepsilon dF_j(\varepsilon) - c_{j0}$ be the expected utility of searching a product net of baseline search costs. Then, the consumer will search the product located at l if its *expected net utility* of searching satisfies

$$\gamma_{j_l} \geq \alpha(l-1). \tag{3}$$

We assume $\gamma_j \geq 0, \forall j \in J$. If $\exists j$ such that $\gamma_j < 0$, then we can restrict our analysis to the subset of J for which $\gamma_j \geq 0$. Products for which this inequality does

[2] If products represent different categories, one can think of ε_j as the highest utility observed from the products considered in that category.

[3] Consistent with the literature, we let $F_j(x) = P(\varepsilon_j < x)$.

not hold are not candidates for search even when search costs are minimal (i.e. when placed in the least costly location, $l = 1$), meaning that the consumer will never find it optimal to search them. Thus, our assumption is without loss of generality in a model of consumer search. Under this assumption, the consumer is always willing to search the product located at $l = 1$.

Retailer's location problem The products the consumer seeks to purchase are manufactured by third-party sellers and sold through a monopolist retailer. As a result, we assume the retailer cannot change any product features (including prices), because these are under the control of each seller. Rather, the retailer chooses whether to prioritize product 1 or 2 in location $l = 1$, while taking into account the higher search cost for the product placed in location $l = 2$. The locations may be physical locations within a brick and mortar store, such as placing the product near the entrance versus in the back or by the checkout area. In the context of an online retailer, different locations may be featuring the product on the homepage, prioritized in a menu, versus accessible only through keyword search. In Section 5, we relax the assumption that the retailer cannot change product features, and consider the case where the retailer chooses both location and product prices, and show that our results continue to hold.

The retailer chooses products to place in each location, $\overrightarrow{j}^* = (j_1^*, j_2^*)$, where j_l^* represents the option displayed in location l, in order to maximize its expected payoff. If the consumer does not search a product in location l, the retailer gets no payoff from that product. If the consumer does search, we assume the retailer gets a payoff of 1 if the consumer purchases the product, and 0 otherwise (Athey and Ellison 2011). The model extends easily to accommodate more general payoffs, such as ones that depend on prices, without changing the solution structure. The case when the retailer chooses prices as well as locations is more complex and is addressed in Section 5. In this setting, the retailer's expected payoff is given by:

$$\overrightarrow{j}^* = \arg\max_{\overrightarrow{j}} \sum_{l=1}^{2} \mathbb{1}[\gamma_{j_l} \geq \alpha(l-1)] \cdot \phi_{j_l}, \tag{4}$$

where the first term is equal to 1 if the consumer searches the product in location l, and 0 otherwise, and the second term gives the probability that the consumer buys the product conditional on search.

4 Optimizing product location

Maximizing the number of products searched Before solving the general product allocation problem where the retailer maximizes its payoff given in Eq. 4, we first consider a simpler objective function, which will help build intuition for the main problem: maximizing the number of products which the consumer searches, that is solving

$$\overrightarrow{j}^* = \arg\max_{\overrightarrow{j}} \sum_{l=1}^{2} \mathbb{1}[\gamma_{j_l} \geq \alpha(l-1)]. \tag{5}$$

Proposition 1 describes where the retailer should allocate products to maximize the number of products searched by the consumer, while taking her optimal search decision into account.

Theorem 1 *With two products* $j \in \{1, 2\}$ *and* $\alpha > 0$*, prioritizing the product with lower expected net utility* γ_j *in location* $l = 1$ *maximizes the number of products searched.*

Proof Depending on the values of γ_j, there are three cases to consider, depicted in Fig. 1 below. The dots in the figure represent γ_j values relative to the value of α. Let $|j : \gamma_j \geq \alpha|$ denote the number of products satisfying condition $\gamma_j \geq \alpha$, and recall that $\gamma_j \geq 0$ for both products.

- Case 1: $|j : \gamma_j \geq \alpha| = 1$
  In this case, the consumer will search the product with $\gamma_j \geq \alpha$ regardless of whether it is located in the high or the low cost location. However, the consumer will only search the other product if it is in the low cost location. Thus, prioritizing the product with lower expected net utility, that is the one with $\gamma_j < \alpha$, by placing it in the low cost location, strictly increases the number of products searched.
- Case 2: $|j : \gamma_j \geq \alpha| = 2$
  In this case, the consumer will search both products regardless of which locations they occupy. Thus, any allocation maximizes the number of products searched, including placing the product with lower expected net utility in the lower cost location.
- Case 3: $|j : \gamma_j \geq \alpha| = 0$
  The consumer will only search the product in the low cost location. Therefore, any allocation, including placing the product with lower expected net utility in the lower cost location, maximizes the number of products searched. □

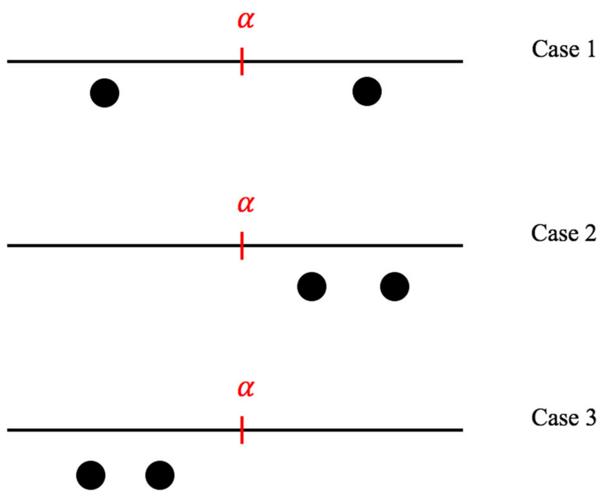

Fig. 1 Illustrating the possible values of γ_j with two products

The intuition for this result is simple: when search costs vary by location, the retailer has an incentive to place the lower utility products in the low cost locations, because the high utility products are more likely to be searched even at higher costs. If instead high utility products occupy the low cost locations, lower utility products might not get searched. Thus, placing lower utility products in the low cost locations maximizes the number of products searched by the consumer.

Maximizing expected payoffs We now turn to the general product allocation problem, where the retailer maximizes expected payoffs given in Eq. 4, using the intuition developed above. Proposition 2 describes our result.

Theorem 2 *With two products $j \in \{1, 2\}$ and $\alpha > 0$, to maximize expected payoffs:*

(i) *if $|j : \gamma_j \geq \alpha| \geq 1$, the retailer will prioritize the product with lower expected net utility γ_j in location $l = 1$;*
(ii) *if $|j : \gamma_j \geq \alpha| = 0$, the retailer will prioritize the product with higher purchase probability ϕ_j in location $l = 1$.*

Proof We consider three cases separately. In the first case, only one product has an expected net utility greater than α, that is $|j : \gamma_j \geq \alpha| = 1$. In the second case, both products have expected net utility greater than α, that is $|j : \gamma_j \geq \alpha| = 2$. The condition $|j : \gamma_j \geq \alpha| \geq 1$ corresponds to these two cases. The third case is when both $\gamma_1, \gamma_2 < \alpha$ or $|j : \gamma_j \geq \alpha| = 0$.

- Case 1: $|j : \gamma_j \geq \alpha| = 1$
  In this case, there exists a product allocation under which the consumer will search both options. As per Proposition 1, this is accomplished by placing the product with $\gamma_j \geq \alpha$ in the higher cost location, $l = 2$, that is $j_2^* = \arg\max_j \gamma_j$, and by placing the remaining product in location $l = 1$. This strategy leads to an expected payoff of $\phi_1 + \phi_2$ for the retailer, as both products get searched, which is greater than the expected payoff if only one product is searched, ϕ_j.
- Case 2: $|j : \gamma_j \geq \alpha| = 2$
  The consumer will search both options. Thus, any allocation maximizes the retailer's expected payoff, including prioritizing the lower expected net utility product.
- Case 3: $|j : \gamma_j \geq \alpha| = 0$
  The consumer only searches one product in this case. Therefore, to maximize its expected payoff, the retailer prioritizes in location $l = 1$ the product with the highest purchase probability, or $j_1^* = \arg\max_j \phi_j$. □

Our results in Proposition 2 show that, for some search cost parameter values, the retailer may choose to position the products in increasing order of utility: lower utility products first, in the easily accessible low search cost locations, and higher utility products later, in the less accessible, harder to search locations. Therefore, the retailer's optimal product location depends not only on product utilities, but also on the exact values of search costs in different locations.

It is worth contrasting our results with those in the information retrieval literature. This literature studies the problem of ordering products faced by search engines that need to present search results to consumers by means of a ranked list. Search engines compete with each other primarily on their ability to place the most relevant search results at the top of the ranking. Thus, information retrieval algorithms are focused on ranking results in decreasing order of utilities. In contrast, retailers are more differentiated from each other, on factors such as physical location, product assortments, customer service, loyalty programs, etc. Therefore, what we have shown here is that retailers do not necessarily order products by utility, and must take the relative search costs into account when choosing product locations within a store. Recall that the result applies for the products that have above 0 expected net utility from search ($\gamma_j \geq 0$).

Also, it is worth noting how our results thus far compare to the benchmark case when $\alpha = 0$, that is when search costs are not a function of a product's location, as assumed by most of the consumer search literature. In this case, since her choices are unaffected by the retailer's product location decision, the consumer will search all available products with $\gamma_j \geq 0$ and will purchase a product if its realized match value is positive. This leads to retailer expected payoffs that equal $\sum_{j \in J} \phi_j$, and to a consumer total expected utility that equals $\sum_{j \in J} \gamma_j$. These payoffs are (weakly) larger than when $\alpha > 0$, since $\alpha = 0$ implies search costs are lower, meaning that the consumer searches more and is more likely to purchase. Thus, the case $\alpha > 0$ makes both the retailer and the consumer (weakly) worse off.

5 Extensions

In this section, we tackle several extensions: (i) the retailer chooses both the product location and each individual product's price; (ii) there are $N \geq 2$ products available, rather than just two; (iii) products are substitutes and the consumer has unit demand; (iv) the focal retailer faces competition. We consider these extensions in sequence.

5.1 The choice of both product location and prices

In the baseline model presented in Section 3, we focused on the case where the retailer solves an optimal product location problem, abstracting away from other issues, such as optimal pricing. In this section, we extend the model by allowing the retailer to choose both a product's location and its price, and the consumer searches to reveal both her match with a product and its price. Our setup will mirror that in Wolinsky (1986) and Anderson and Renault (1999), except that the consumer will search among two independent products, and the retailer will choose not only prices, but also product locations.

Suppose a consumer searches for both match values ε_j and prices p_j, where $j \in \{1, 2\}$. As before, the retailer cannot observe ε_j, but it is common knowledge that match values are distributed according to $F_j(\cdot)$ defined on the interval $\left[\underline{\theta}_j, \overline{\theta}_j\right]$.

The retailer sets prices and the consumer expects to see the equilibrium price p_j^* before searching j. Also, the retailer chooses a product's location, which affects the consumer's search cost. Formally, the retailer chooses products to place in each location, $\overrightarrow{j}^* = (j_1^*, j_2^*)$ and prices $\overrightarrow{p}^* = (p_{j_1^*}^*, p_{j_2^*}^*)$, where j_l^* represents the product in location l and $p_{j_l^*}^*$ represents its equilibrium price. In this setting, the consumer makes optimal search decisions by choosing to search product j_l^* if and only if

$$\int_{p_{j_l^*}^*}^{\overline{\theta}_{j_l^*}} (\varepsilon - p_{j_l^*}^*) dF_{j_l^*}(\varepsilon) - c_{j_l^*} \geq 0, \tag{6}$$

and will purchase product j_l^* if the difference between her realized match value and the observed price (which may differ from the equilibrium price) is greater than the outside option, that is if $\varepsilon_{j_l^*} - p_{j_l^*} \geq 0$.

The retailer maximizes expected payoffs, which are given by

$$\max_{\overrightarrow{j}, \overrightarrow{p}} \sum_{l=1}^{2} \mathbb{1}[\gamma(p_{j_l}) \geq \alpha(l-1)] \cdot [1 - F_{j_l}(p_{j_l})] \cdot p_{j_l}, \tag{7}$$

where $\gamma(p_j) = \int_{p_j}^{\theta_j} (\varepsilon - p_j) dF_j(\varepsilon) - c_{j0}$.

Equilibrium To maximize expected payoffs, the retailer would set equilibrium prices according to the first order condition in Eq. 7, that is it would set $p_{j_l^*}^* = \frac{1-F_{j_l^*}(p_{j_l^*}^*)}{f_{j_l^*}(p_{j_l^*}^*)}$, if at that price the consumer would search j_l^*. However, depending on the value of α, at price $p_{j_l^*}^*$, the consumer might not be willing to search the product located at l, that is $\gamma(p_{j_l^*}^*) \geq \alpha(l-1)$ might not hold. In this case, the retailer maximizes expected payoffs by solving for the price that makes the consumer indifferent between searching and not searching the product located at l, that is by solving for p from $\gamma(p) = \alpha(l-1)$, as long as the resulting price is non-negative; otherwise, the retailer will prefer not to offer the product. Thus, equilibrium prices depend on the value of α.

To maximize expected payoffs, the retailer chooses not only equilibrium prices, but also each product's location. In this case, it faces a tradeoff. On the one hand, placing the product with relatively higher expected net utility in the lower cost location $l = 1$, means being able to charge a high price for it (both because of its higher utility and because of the lower search cost), but risking that the consumer will not search the other product (both because of its lower utility and because of the higher search cost in location 2). On the other hand, placing the product with relatively lower expected net utility in the lower cost location $l = 1$ might involve charging a lower price than in location $l = 2$ for it, but ensuring that both products will be searched. Depending on the cost (α) of searching the product in location 2, equilibrium prices and the optimal product allocation will change, and one of these strategies will be more profitable.

Analyzing jointly the optimal product location and the optimal pricing decision of the retailer is complex. To make progress, we make the following two assumptions:

(i) the two products, called L and H, have match values that are uniformly distributed according to $U(0, L)$ and $U(0, H)$, where $H > L > 0$, and (ii) baseline search costs are zero for both products. Given our distributional assumptions, the expected net utility from searching product $\theta = \{L, H\}$ when it is placed in location l and the consumer expects to observe a price $p^*_{\theta_l}$ for it before search, equals $\gamma(p^*_{\theta_l}) = (\theta - p^*_{\theta_l})^2/(2\theta)$. If the consumer were to search the product θ in location l when it is priced at p_{θ_l}, then the retailer's expected payoff from selling this product would equal $p_{\theta_l}(\theta - p_{\theta_l})/\theta$. In what follows, we study where the retailer should locate L and H and what it should charge each, taking into account how these choices will influence the consumer's optimal search decision.

Consider first the case when the retailer places product L in location $l = 1$. The retailer will charge $p^*_L = L/2$ (the price that solves $p = \frac{1-F_L(p)}{f_L(p)}$), because at this price the consumer is willing to search L, that is the expected net utility from searching $(L - p^*_L)^2/(2L) = L/8$ is non-negative. Then, depending on the value of α, the price of the other product, H, will vary. In particular, for $\alpha \leq H/8$, charging $p_H = H/2$ is optimal, because at this price, the expected net utility from searching H, which equals $(H - p^*)^2/(2H)$, (weakly) exceeds the additional cost, α. As α increases, the retailer will charge a lower price for H to ensure that the consumer will search it given the larger cost of searching a product in location $l = 2$. More precisely, for $H/8 < \alpha \leq H/2$, the price of H equals $p^*_H = H - \sqrt{2H\alpha}$ (obtained by solving for p from $(H - p)^2/(2H) = \alpha$), which decreases in α. At some point, when $\alpha > H/2$, because there exists no non-negative price at which the consumer would search H, the retailer will prefer not to offer it. We illustrate this pricing pattern in Fig. 2a below, assuming $L = 1$ and $H = 2$.

The expected payoffs of the retailer when it places product L in location $l = 1$, π_L, are (weakly) decreasing in α, as illustrated in Fig. 2b. More precisely, when $\alpha \leq H/8$, expected payoffs are constant at $\pi_L = (H + L)/4$. For $H/8 < \alpha \leq H/2$, expected payoffs equal $\pi_L = L/4 + \sqrt{2H\alpha} - 2\alpha$, decreasing in α over the range. When $\alpha > H/2$, the price of H that would make the consumer willing to search H in location $l = 2$ is no longer positive, so the retailer will prefer not to offer H, and

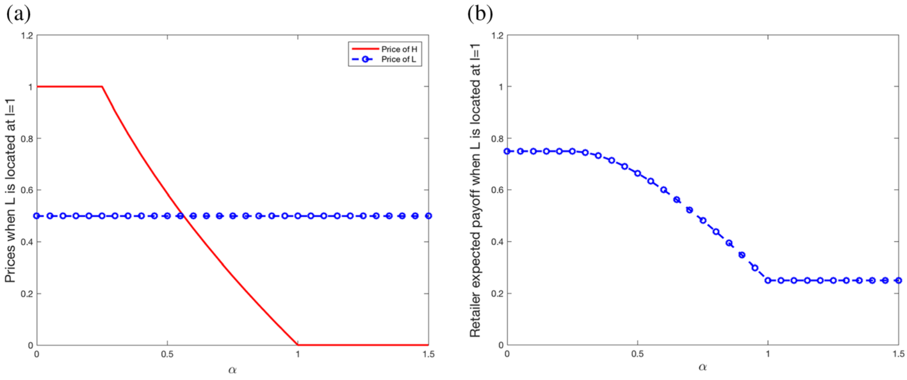

Fig. 2 Product prices and retailer expected payoff when the lower expected utility product is placed in the low cost location

expected payoffs are constant as a function only of the price of L, equal to $L/4$. To summarize, the retailer's expected payoffs when it places L in location $l = 1$ are given by

- If $\alpha \leq H/8$, then $\pi_L = \frac{H+L}{4}$ and retailer offers both products.
- If $H/8 < \alpha \leq H/2$, then $\pi_L = \frac{L}{4} + \sqrt{2H\alpha} - 2\alpha$ and retailer offers both products.
- If $\alpha > H/2$, then $\pi_L = \frac{L}{4}$ and retailer only offers product L.

A similar analysis can be performed for the case when the retailer places H in location $l = 1$. We illustrate the results of this analysis as a function of α in Fig. 3. In sum, the retailer's expected payoffs when H is located at $l = 1$, π_H, are given by

- If $\alpha \leq L/8$, then $\pi_H = \frac{H+L}{4}$ and retailer offers both products.
- If $L/8 < \alpha \leq L/2$, then $\pi_H = \frac{H}{4} + \sqrt{2L\alpha} - 2\alpha$ and retailer offers both products.
- If $\alpha > L/2$, then $\pi_H = \frac{H}{4}$ and retailer only offers product H.

Notice that, as in the case where the retailer chooses only the product location (not prices), the retailer is (weakly) worse off when $\alpha > 0$ rather than when search costs are not a function of a product's location, that is when $\alpha = 0$. In addition, as α increases, the retailer's expected payoff (weakly) decreases.

Given these results, we can now determine the equilibrium product allocation and each product's price. Because $L < H$, the retailer will find it profitable to stop offering a second product sooner (i.e., for smaller values of α) if it places H rather than L in location $l = 1$. Also, when α is small, the retailer is indifferent between placing L or H in location $l = 1$. These two facts mean that there are three payoff relevant ranges of α for the retailer, as illustrated in Fig. 4. First, for relatively low values of α, when $\alpha \leq L/8$, placing either L or H in location $l = 1$ leads to the same expected payoffs for the retailer, and the consumer is willing to search both products. Second, for relatively high values of α (obtained by solving for α when $\pi_H = H/4$ and $\pi_L = L/4 + \sqrt{2H\alpha} - 2\alpha$), the retailer will prefer to place H in location $l = 1$, making higher payoffs although the consumer would only search one

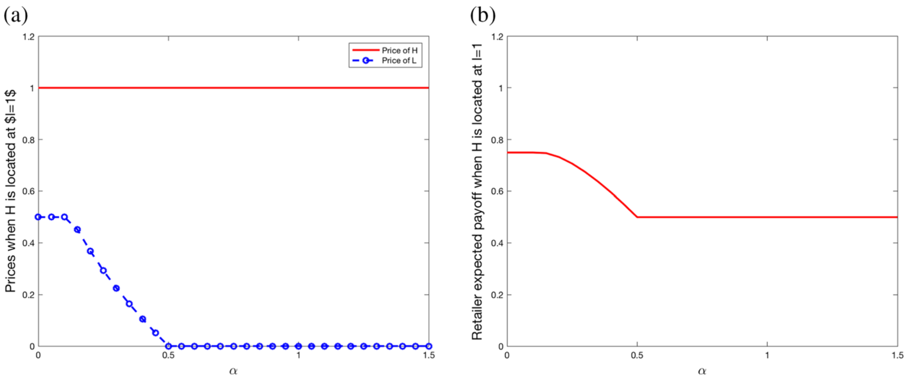

Fig. 3 Product prices and retailer expected payoff when the higher expected utility product is placed in the low cost location

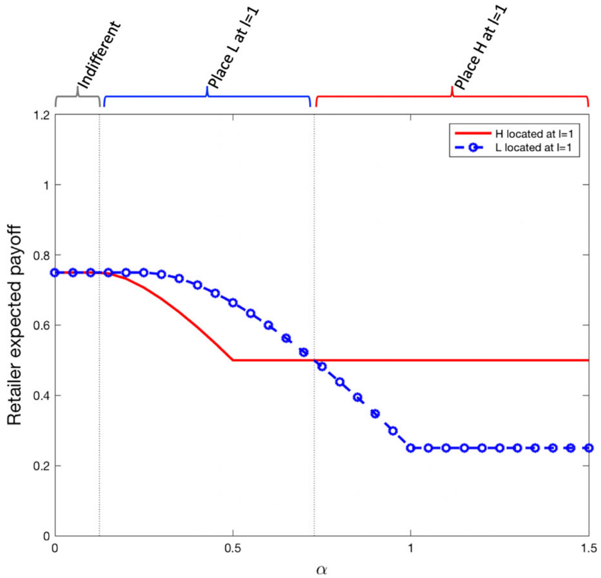

Fig. 4 Retailer expected payoff and the optimal product allocation as a function of α

product. Finally, for intermediate values of α, the retailer will prefer to place L in location $l = 1$, and the consumer will search both products.

The result in Fig. 4 shows that when α is not too large to prevent the consumer from searching the product in location $l = 2$, the retailer will (at least weakly) prefer to place L in location $l = 1$. This encourages the consumer to search both options, and it maximizes the retailer's expected payoffs. Only when searching the product in location $l = 2$ is prohibitively expensive, then the retailer prefers to place H in location $l = 1$, and the consumer will only search one product. This mirrors our result in Proposition 2 when the retailer chose only the product location. Thus, our main insight from the baseline model continues to hold when the retailer chooses both product location and product prices.

In addition, we are able to determine the equilibrium prices that would prevail as a function of the product allocation chosen. Figure 5 illustrates our result by combining Figs. 2a and 3a. As can be seen, for most values of α, regardless of the product allocation chosen, the retailer charges a higher price for H, than for L. The only exception is in the area between the two bars where the retailer places L in location $l = 1$ and the consumer searches both products. In this case, because of the relatively high value of α, the retailer needs to decrease the price of H sufficiently to encourage the consumer to search it. In sum, we find that α, the cost of searching a product when it is located in $l = 2$, affects not only the retailer's product location decision, but also its optimal pricing decision. Thus, taking into account how product location affects consumer search costs, is of even greater importance when the retailer chooses product prices in addition to the locations of the products available within the store.

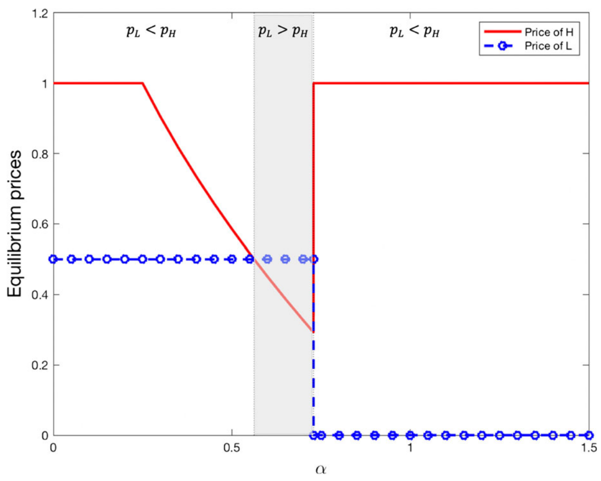

Fig. 5 Equilibrium prices as a function of α

5.2 Optimal product location with $N \geq 2$ products

The baseline model focused on the problem of optimally allocating two products within a store. We now consider the optimal allocation algorithm for any number of products $j \in J = \{1, \ldots, N\}$, where $N \geq 2$. We assume there are more locations $l \in L = \{1, \ldots, M\}$ available than products ($M \geq N$), that is the retailer is not capacity constrained; otherwise, our results below should be understood as describing the optimal product allocation for the M products with the highest ϕ_j. The following notation will be useful. Define L_j as the location of product j that satisfies

$$\alpha(L_j - 1) \leq \gamma_j < \alpha L_j, \tag{8}$$

that is, L_j is the location with the highest search cost in which the consumer would be willing to search product j, which we refer to as j's *search cutoff*. Also, let

$$R_l = \{j \in J | L_j \geq l\} \tag{9}$$

be the set of products that the consumer would optimally search if located at l. Let S be the set of products resulting from an algorithm and $\overrightarrow{j}^S = \{j_1^S, \ldots, j_{N_S}^S\}$ be the resulting allocation of products, where j_l^S denotes the product placed in location l and N_S gives the total number of locations where products are placed.

Note that the retailer would only display a product if the consumer is willing to search it. The consumer is not willing to search a product if the cost of search is higher than the benefit. Also, after searching a product, all uncertainty about it is revealed, so the consumer would not be willing to search the same product more than once.

We capture these restrictions of optimal consumer search behavior on the retailer's allocation algorithm by defining a *feasible* allocation of products as follows.

Definition 1 Let $\overrightarrow{j}^S = \{j_1^S, \ldots, j_{N_S}^S\}$ be a feasible allocation of products if

1. each allocated product is a candidate for search in the suggested location, that is $L_{j_l^S} \geq l$ for $l \in \{1, \ldots, N_S\}$;
2. each product occupies only one location, that is if $l \neq l'$, then $j_l^S \neq j_{l'}^S$.

Denote by S_l the subset of products placed in lower cost locations than l, and by $\bar{S}_l$ its complement. If the product allocation is feasible for locations of lower cost than l, note that the consumer will optimally choose to search all $j \in S_l$.

Maximizing the number of products searched We will first derive the algorithm that maximizes the number of products searched, that is determine how to locate products to solve

$$\overrightarrow{j}^* = \arg\max_{\overrightarrow{j}} \sum_{l=1}^{N} \mathbb{1}[\gamma_{j_l} \geq \alpha(l-1)]. \tag{10}$$

We will show that, just like in the case of two products, prioritizing options with lower expected net utility maximizes the number of products searched by the consumer.

Consider the following algorithm that results in a feasible set of product allocations, S.

Algorithm 1

Consider locations in increasing order of their search cost. In location l,
1. identify the set of products that have not been placed in a location yet, $\bar{S}_l$;
2. construct the set of products that are eligible to search, R_l;
3. place at l the product that satisfies

$$j_l^S = \arg\min_{j \in \{\bar{S}_l \cap R_l\}} L_j. \tag{11}$$

In other words, Algorithm 1 places a product in location l that satisfies three conditions. First, considering locations starting with the lower cost one, the product located at l must be one that has not yet been been allocated, that is $j_l^S \in \bar{S}_l$, because S is feasible. Second, it must satisfy $j_l^S \in R_l$. Otherwise, the consumer would not be willing to search it. Finally, among the set of products that satisfy both conditions, the algorithm places the product with minimum L_j at l. Note that the resulting product allocation may not be unique, because there might be several options in the set $\bar{S}_l \cap R_l$ at a given l with the minimum search cutoff. In this case, the algorithm will break ties randomly.

The next result shows why we introduced Algorithm 1.

Theorem 3 *When $\alpha > 0$, allocating products to locations as per Algorithm 1 maximizes the number of products searched by the consumer.*

Proof See Appendix. □

Algorithm 1, that maximizes the number of products searched by the consumer, prioritizes options with lower L_j. This means that the resulting product allocation prioritizes options with lower expected net utility, γ_j, since by definition products with lower L_j have lower γ_j and vice versa. Thus, our result in Proposition 3 extends that in Proposition 1 for the case of two products to any number of products.

Example 1 We illustrate our result using the following example. Suppose there are five products, $j \in J = \{1, \ldots, 5\}$ with search cutoffs L_j, as indicated in Table 1 below. Consider locations in increasing order of their search cost. In location $l = 1$, $\bar{S}_1 \cap R_1 = \{1, 2, 3, 4, 5\}$ and Algorithm 1 would place the product with the smallest search cutoff, which is product 1. In location $l = 2$, $\bar{S}_2 \cap R_2 = \{2, 3, 4, 5\}$, and there are two products (2 and 3) with the same minimum search cutoff. Breaking ties randomly, Algorithm 1 would place product 2 in location $l = 2$. In location $l = 3$, the consumer would not consider searching either product 2 or 3, so $\bar{S}_3 \cap R_3 = \{4, 5\}$, and Algorithm 1 would place product 4 because it has the smallest search cutoff, and in location $l = 4$, it would place product 5. Thus, the allocation (not unique) that leads to the maximum number of products searched (which is 4) is given by: [1, 2, 4, 5]. It is straightforward to check that no other product allocation leads to a strictly higher number of products searched.

Maximizing expected payoffs In Proposition 3 above, we showed that prioritizing lower utility products maximizes the number of products searched by the consumer. We now turn to the general problem of maximizing the retailer's expected payoff. The retailer chooses product locations in order to maximize its expected payoff by solving

$$\vec{j}^{\,*} = \arg\max_{\vec{j}} \sum_{l=1}^{N} \mathbb{1}[\gamma_{j_l} \geq \alpha(l-1)] \cdot \phi_{j_l}, \tag{12}$$

where the first term equals 1 if the consumer searches the product in location l, and 0 otherwise, and the second term gives the probability that the consumer buys the product conditional on search. In this setup, we show that the optimal product allocation still prioritizes products with lower utility.

Table 1 Search cutoffs for *Example 1*

j	1	2	3	4	5
L_j	1	2	2	3	5

The problem of maximizing expected payoffs can be mapped into a version of the so called "job sequencing with deadlines" or the "task scheduling" problem in computer science, that deals with scheduling jobs to be processed by a machine (Cormen et al. 2009). Although our setup differs (e.g. the consumer in our model can choose to search products in a different order than the one in which the retailer locates products in the store, while the machine processes jobs in sequence), the solution to our problem can be adapted from the greedy algorithm solution to the task scheduling problem, as we show in Algorithm 2 below.

Consider the following algorithm that produces a feasible set of allocated products, S.

Algorithm 2

1. Sort products $j \in J$ in (weakly) decreasing order of the purchase probability, ϕ_j.
2. Place the product with the highest purchase probability at its search cutoff, L_j.
3. Take the next highest purchase probability product, and place it at its search cutoff location L_j if available; otherwise place it in the most costly available location with a cost lower than its search cutoff. If all remaining unoccupied locations have search costs higher than L_j, then skip product j.
4. Repeat step 3 for every remaining product.

Note, as in the case of Algorithm 1, that the resulting product allocation may not be unique, because several options might have the same purchase probability and search cutoff. In this case, we break ties randomly. Also, note that the resulting product allocation may contain gaps, that is locations where no product is placed, equivalent to the retailer removing shelves from the store. This algorithm maximizes the retailer's expected payoff.

Theorem 4 *When $\alpha > 0$, allocating products to locations as per Algorithm 2 maximizes the retailer's expected payoff.*

Proof See Appendix. □

Example 2 We illustrate this result using the following example. Suppose there are five products, $j \in J = \{1, \ldots, 5\}$, with probability of purchase ϕ_j and search cutoffs L_j, as indicated in Table 2. We sorted options by their purchase probability ϕ_j. Now we can use Algorithm 2 to characterize the optimal product allocation that maximizes retailer expected payoffs. We start with product 1, because it has the highest purchase

Table 2 Purchase probabilities and search cutoffs for *Example 2*

j	1	2	3	4	5
ϕ_j	0.20	0.15	0.10	0.05	0.01
L_j	5	3	1	5	2

Table 3 The final product allocation based on Algorithm 2 for *Example 2*

Location	1	2	3	4	5
j	3	5	2	4	1
ϕ_j	0.10	0.01	0.15	0.05	0.20

probability, $\phi_1 = 0.20$. Given that the consumer would be willing to search this product even when placed in the higher cost location (its search cutoff is $L_1 = 5$), we place it in location 5. We then consider product 2 with $\phi_2 = 0.15$. Because the consumer would search this product in all of the first three locations, and no other product is placed at its search cutoff, we place product 2 in location 3. Similarly, we place product 3 in location 1. Product 4 would be searched by the consumer in all of the first five locations, but since location 5 is occupied, we place it in the highest search cost location available, location 4. Finally, we place product 5 at its search cutoff, location 2, since it is available. The final product allocation is [3,5,2,4,1], leading to expected total payoffs of 0.51. We illustrate this final product allocation in Table 3. It is straightforward to check that no other allocation will lead to a strictly higher payoff.

It is worth noting that in the product allocation that maximizes the retailer's expected payoffs, options are not ordered by purchase probabilities, ϕ_j. For example, the product with the highest purchase probability (product 1) is placed in the last location. Furthermore, any attempt to place the product with the highest purchase probability in the first location would decrease the retailer's expected payoff.

Rather, as we demonstrate below (see Proposition 5), products are placed in locations such that those with lower search cutoffs L_j are prioritized in lower search cost locations. For instance, in the example we just discussed, in the final product allocation, the search cutoffs of the products allocated are [1,2,3,5,5]. By definition, products with a lower search cutoff L_j also have a lower expected net utility γ_j. Thus, just like in the case of two products, prioritizing options with lower expected net utility maximizes the retailer's expected payoff. We now show more formally that in the optimal product allocation, products with lower expected net utility are prioritized among those allocated. We focus on the case where the consumer is willing to search at least two products in the optimal product allocation, because if not, it is futile to think about comparing products on any criterion.

Theorem 5 *Let S be the set and $\overrightarrow{j}^S$ be the product allocation resulting from Algorithm 2 that maximizes retailer expected payoffs. Suppose the consumer is willing to search at least two products under this optimal product allocation. Then $\overrightarrow{j}^S$ prioritizes products with lower expected net utility, γ_j, or can be rearranged in such a way while preserving the same maximum expected payoff.*

Proof See Appendix. □

In sum, what we have shown is that when the retailer chosen product locations affect consumer search costs, prioritizing products with lower utility, as per Algorithm 2, maximizes the retailer's expected payoff. This finding demonstrates that our main result for the case of two products in Proposition 2, extends to any number of products.

5.3 Product substitutes

In this section, we extend our model and consider the case where products are substitutes, rather than independent. This setting is the most common one found in the theoretical search literature (e.g. Weitzman, 1979), modeling the case where the consumer is searching for the best alternative within a product category, making at most one purchase (i.e. unit demand). For this setting, we derive the optimal product allocation that maximizes the retailer's expected payoff, while taking into account how the allocation created affects the consumer's optimal search decision.

Consider a retailer selling two products $j \in \{1, 2\}$. Products are substitutes and the consumer seeks to purchase at most one product, or choose the outside option of not purchasing with expected utility normalized to zero. Searching a product j reveals all uncertainty about it, giving the consumer a potential match utility ε_j assumed to be drawn from a product specific distribution function $F_j(\cdot)$ defined on the interval $[\underline{\theta}_j, \overline{\theta}_j]$. Both the consumer and the retailer know the distribution of match values.

The payoff of the retailer from a purchase made by the consumer equals 1. As before, the retailer chooses whether to prioritize product 1 or 2 in location $l = 1$, taking into account the higher search costs of the product located at $l = 2$. We continue to consider only products for which $\gamma_j \geq 0$, because those with negative expected net utility are not candidates for search even when search costs are minimal (i.e. when placed in location $l = 1$). Under this assumption, the consumer is willing to search the product in location $l = 1$.

Unlike the case considered in the baseline model where products are independent, when products are substitutes, what the consumer observes while searching one product affects her decision to search another product, to stop searching, and to make a purchase decision. Thus, her optimal search decision depends additionally on which products she has already searched. Before searching any product, the consumer's best utility observed so far is the outside option with expected utility normalized to zero, so the consumer will search the product in location l if

$$\int_{0}^{\overline{\theta}_{j_l}} \varepsilon dF_{j_l}(\varepsilon) - c_{j_l} \geq 0. \tag{13}$$

After searching product j_l, the consumer observes a utility ε_{j_l}, and will search product $k_{l'}$ in location l' if

$$\int_{z_{j_l}}^{\overline{\theta}_{k_{l'}}} (\varepsilon - z_{j_l}) dF_{k_{l'}}(\varepsilon) - c_{k_{l'}} \geq 0, \tag{14}$$

where $z_{j_l} = \max\{0, \varepsilon_{j_l}\}$ gives the best utility observed so far.

In this setting, we now consider the optimal product allocation when the retailer maximizes its expected payoff. Proposition 6 below shows our result for the case of product substitutes. This result is the equivalent of the one in Proposition 2, that dealt with independent products.

Theorem 6 *With two product substitutes $j \in \{1, 2\}$ and $\alpha > 0$, to maximize expected payoffs*

(i) *if $|j : \gamma_j \geq \alpha| \geq 1$, the retailer will prioritize the product with lower expected net utility γ_j in location $l = 1$;*
(ii) *if $|j : \gamma_j \geq \alpha| = 0$, the retailer will prioritize the product with higher purchase probability ϕ_j in location $l = 1$.*

Proof If $|j : \gamma_j \geq \alpha| \geq 1$, then there are two cases to consider: either $|j : \gamma_j \geq \alpha| = 1$ or $|j : \gamma_j \geq \alpha| = 2$. The case $|j : \gamma_j \geq \alpha| = 0$ corresponds to the scenario in which both products have expected net utility lower than α. We now consider each case separately, to prove the claim.

- Case 1: $|j : \gamma_j \geq \alpha| = 1$
  Suppose, without loss of generality, that $\gamma_1 \geq \alpha, \gamma_2 < \alpha$. If the retailer places product 1 in location $l = 1$, the consumer searches it and reveals a match value of ε_1, which leads to a maximum utility observed so far of $z_1 = \max\{0, \varepsilon_1\}$. The consumer will then not search product 2, because $\gamma_2 = \int_0^{\bar{\theta}_2} \varepsilon dF_2(\varepsilon) - c_{20} < \alpha$, meaning that $\int_{z_1}^{\bar{\theta}_2} (\varepsilon - z_1) dF_2(\varepsilon) - c_{20} < \alpha$. Thus, the retailer's expected payoff equals $\phi_1 = 1 - F_1(0)$.
  Suppose instead that the retailer places product 2 in location $l = 1$. Now both products are eligible to be searched when the best utility observed up to that point is zero. Thus, if the first product searched reveals a negative reward, the consumer will search the remaining product and buy if $\varepsilon_j \geq 0$. Otherwise, if the first searched product reveals a positive reward, the consumer will certainly buy (regardless of whether she searches the other product or not). This means that the consumer will only not purchase if both products reveal a negative reward, making the expected payoff of the retailer equal to $1 - F_1(0)F_2(0)$. Because this value is greater than $1 - F_1(0)$, that is the expected payoff from placing product 1 in location $l = 1$, the retailer will prefer to prioritize the product with lower expected net utility (product 2 in this case) in location $l = 1$.
- Case 2: $|j : \gamma_j \geq \alpha| = 2$
  In this case, regardless of the product allocation, both products are eligible to be searched when the best utility observed up to that point is zero. Using similar logic as in case 1 above, the consumer will only not purchase if both products reveal a negative reward, making the expected payoff of the retailer equal to $1 - F_1(0)F_2(0)$. Thus, prioritizing the product with lower expected net utility in location $l = 1$ maximizes the retailer's expected payoff.
- Case 3: $|j : \gamma_j \geq \alpha| = 0$

The consumer will only search one product, regardless of the product allocation. Thus, prioritizing the product with the highest purchase probability in location $l = 1$ maximizes the retailer's expected sales. □

In sum, we find that, as in the case of independent products, when products are substitutes, the retailer has an incentive to prioritize the product with lower expected utility in location $l = 1$. This result reinforces the importance of accounting for the impact of product location on consumer search costs when studying how to optimally locate products.

5.4 Retailer competition

So far, we have focused on the optimal product location problem of a single retailer. In this section, we study how the presence of competition affects our results. In particular, we show that when retailers compete solely on their product allocation, then our results do not hold anymore and prioritizing higher utility products is optimal. This mirrors the result obtained by algorithms designed for search engine information retrieval. However, in reality retailers may be differentiated in several dimensions, therefore not competing directly on their product allocation. For example, retailers may differ in terms of their physical location (e.g. one of the few grocery or department stores in a given neighborhood), allowing them to enjoy a certain degree of market power since transportation costs may prevent consumers from visiting multiple retailers. Also, retailers may differ in their product assortment, as is the case for grocery stores such as Whole Foods or Safeway. Or retailers may provide value to consumers through customer service or loyalty programs that make it costly for them to switch retailers. Finally, retailers may compete on all these and other dimensions at the same time. When this is the case, we show that the optimal product location solution is still to prioritize the lower utility products in the lower search cost locations.

Suppose there are two retailers $\{A, B\}$ that differ in their product allocation, that is they choose to place products differently in the store. The consumer chooses not only which product to search/purchase after visiting a retailer, but also which retailer to visit. We assume the cost to visit both retailers exceeds the benefit, so the consumer will only visit one retailer, allowing us to identify which one she prefers.

In this setting, we consider two extreme cases. First, suppose there are two products $j \in \{1, 2\}$ available, with $\gamma_1 > \gamma_2 \geq 0$, that both retailers sell them, and that they differ only in their product allocation: retailer A places product 2 in location $l = 1$, while retailer B places product 1 in location $l = 1$. So A allocates products as per our solution, while B allocates the higher expected net utility product in location $l = 1$. Then, retailer B will be (weakly) preferred by the consumer. To see this, consider again all possible cases determining the relation between the expected net utility and α. For example, if $|j : \gamma_j \geq \alpha| = 1$ and, without loss of generality, $\gamma_1 \geq \alpha$ and $\gamma_2 < \alpha$, then by visiting A the consumer's expected benefit equals $\gamma_1 + \gamma_2 - \alpha$, which is less than her benefit from visiting B, γ_1, so the consumer prefers retailer B. Also, if $|j : \gamma_j \geq \alpha| = 2$, the consumer is indifferent between the two retailers.

Finally, if $|j : \gamma_j \geq \alpha| = 0$, then by visiting A the consumer expects a benefit of γ_2, which is lower than her benefit from visiting B, γ_1.

Second, suppose retailers differ not only in their product allocation, but also in their product assortment. For example, suppose that there are four products $j \in \{1, 2, 3, 4\}$ available, with $\gamma_1 > \gamma_2 > \gamma_3 > \gamma_4 \geq 0$, and that A sells products $\{1, 2\}$, and places 2 in location $l = 1$, while B sells products $\{3, 4\}$ and places 3 in location $l = 1$. In other words, A allocates products as per our solution, while B places the higher expected net utility product in location $l = 1$. There are several cases determining the relation between the expected net utility and α, but the relevant ones involve how γ_1 and γ_4 relate to α. We analyze all possibilities below.

- If both $\gamma_1 \geq \alpha$ and $\gamma_4 \geq \alpha$ (implying that $\gamma_j \geq \alpha, \forall j \in \{1, 2, 3, 4\}$), then the consumer prefers to visit A, expecting a benefit of $\gamma_1 + \gamma_2 - \alpha$, to B, where she would obtain only $\gamma_3 + \gamma_4 - \alpha$.
- If $\gamma_1 \geq \alpha$, but $\gamma_4 < \alpha$, then regardless of how γ_2 and γ_3 relate to α, since $\gamma_2 > \gamma_3$, the consumer prefers to visit A, since she expects a benefit of $\gamma_1 + \gamma_2 - \alpha$, which is greater than the benefit of γ_3 from visiting B.
- If both $\gamma_1 < \alpha$ and $\gamma_4 < \alpha$ (implying that $\gamma_j < \alpha, \forall j \in \{1, 2, 3, 4\}$), then the consumer prefers to visit A, where the expected benefit equals γ_2, to B, where her expected benefit would only be γ_3, since $\gamma_2 > \gamma_3$.

As just analyzed, regardless of the relation between the expected net utility and α, the consumer prefers visiting retailer A to B, even though A prioritizes the lower expected net utility product in location $l = 1$. This happens because A's product assortment is superior to that of B (i.e. $\gamma_2 > \gamma_3$).

In sum, we expect that when retailers enjoy the high degree of differentiation observed in many real life examples, our results will continue to hold even in the presence of competition.

6 Discussion

This article addresses the problem of how a multi-category retailer should optimally locate products within a store, when the locations chosen affect consumer search costs. We show that the retailer has an incentive to prioritize products with lower utility. This result is robust to various modifications of the baseline model and continues to hold when the retailer chooses not only product locations, but also prices, when products are substitutes rather than independent, as well as when the focal retailer faces competition. Nevertheless, few issues deserve further study.

Our model considers the optimal location problem of a retailer for a given consumer, that is we focus on the case where consumers visiting the store have similar preferences and search costs (e.g. representative consumer), or where the retailer can adjust the allocation of products to each visiting consumer (e.g. by personalizing the e-commerce website). In reality, heterogeneous consumers with different match value distributions and search costs, may simultaneously enter a store and have to be accommodated by a single product allocation. In these cases, issues such as the percent of each type of consumer in the store and the value of promoting niche products, will determine the optimal allocation of products. We leave these issues to future research.

Also, in our model, the retailer knows both the distribution of match values and consumers' search costs. These assumptions are common in the literature, and are meant to represent the large amounts of data retailers nowadays have access to on consumers searching in their stores. However, in absence of such data, the optimal product location solution might differ from the one we presented in this article.

Our paper does not study how a retailer would allocate complementary products. Instead, we focus on independent and substitute products. There are several reasons for this choice. First, we are not aware of papers describing how consumers would search for complementary products. Rather, complementary products are defined in terms of joint purchase decisions. Second, for certain representations of consumer search for complementary products, our results continue to hold. More precisely, suppose the consumer searches for two complementary products, and will only buy if she finds good matches for both. One possible way to model search for complementary products is then to assume that the observed utility at the first searched product is correlated with the probability of searching its complement. For example, observing a good match when searching for pasta increases the probability that the consumer searches for tomato sauce as well, while a poor match decreases this probability. In this case, the consumer will only buy if she searches both products, and the probability of searching both products is maximized by our solution.

It is worth noting that a special case of our model is one where a consumer walks through the store in a particular order, deciding whether or not to examine encountered product categories, and the retailer decides the order in which the consumer is exposed to the different product categories. For example, stores of the furniture retailer IKEA are often designed in a way that guides customers through a particular path through the store.[4] As the consumer evaluates more products, her search costs may increase due to fatigue, as shown in the consumer behavior literature (Redden 2008) and empirical search literature (Koulayev 2014). Thus, the later within a shopping trip the consumer is exposed to a category, the higher will be the cost of searching that category. Prior work showed that the presence of increasing search costs means that changing the order in which product options are presented can affect the final option chosen by the consumer (Levav et al. 2010) and that search fatigue can influence firm's pricing decisions in equilibrium (Carlin and Ederer 2018). By thinking of the retailer's product order decision in the presence of search fatigue as choosing which products will have higher search costs, one observes the similarity of this problem and the product allocation problem studied in our paper.

Finally, our results can be connected to recent reports that online companies, such as Facebook and Netflix, may prioritize less valuable content to consumers. For example, Facebook orders the sizable amount of content created by a user's friends in her news feed such that less preferred items are ranked at the top (among those the

[4] A heat map describing navigation patterns at Ikea and showing that consumers' path through the store mirrors the store plan can be found at http://www.dailymail.co.uk/femail/article-1349831/Ikea-design-stores-mazes-stop-shoppers-leaving-end-buying-more.html.

user is interested in reading).[5] On Netflix, in order to continue watching a show a user started, she will first have to scroll through other categories, such as "featured content", "Netflix originals", or "trending now" shows, which she is likely less interested in watching next than the show she is currently in the middle of watching.[6] By prioritizing lower utility options, Facebook and Netflix encourage consumers to search more, which in turn benefits them as our proposed strategy suggests (in the case of Facebook through more time spent on the site which may attract more advertisers, and in the case of Netflix through product discovery which might serve to increase long term revenues for the company).

Acknowledgments We are thankful for comments from Alixandra Barasch, Kristina Brecko, Xinyu Cao, Pradeep Chintagunta, Babur De los Santos, Chaim Fershtman, Tobias Gamp, Konstantin Korotkiy, Song Lin, Dmitry Lubensky, Eitan Muller, Cem Ozturk, Vaiva Petrikaite, Robbie Sanders, Andrey Simonov, Adam Smith, Monic Sun, Artem Timoshenko, Miguel Villas-Boas, Chris Wilson, Hema Yoganarasimhan, and attendees of the 2018 Consumer Search and Switching Cost Workshop, the 2018 Workshop on Multi-Armed Bandits and Learning Algorithms, and the 2018 Marketing Science conference. The usual disclaimer applies.

Appendix

Before proving Proposition 3 that Algorithm 1 maximizes the number of products searched by the consumer, we first demonstrate Lemma 1 below, which requires the following notation. Consider locations in increasing order of their search cost. Let $C_l(\{L_j, j \in \overline{S_l}\})$ be the maximum number of products that can be searched starting with location l, given the set of products not yet allocated $\overline{S_l}$ and their search cutoffs L_j.

Lemma 1 *For any* l, $C_l(\{L_j, j \in \overline{S_l}\}) \geq C_l(L_j, j \in \overline{S_l} \setminus \{a\} \cup \{b\})$ *if* $L_a \geq L_b$.

In other words, swapping a yet-to-be-allocated product for a different product with weakly smaller search cutoff cannot increase the total number of products searched.

Proof of Lemma 1 Any allocation of products that maximizes the number of products searched from the set $\overline{S_l} \setminus \{a\} \cup \{b\}$ can also be used to generate an equal number of products searched from $\overline{S_l}$, if it does not place product b in any location. If the search maximizing product allocation does place product b in a location, it can be swapped for product a, because $L_a \geq L_b$. □

Using the result in Lemma 1, we can now prove Proposition 3.

Proof of Theorem 3 Let S and Q be the sets of products resulting from Algorithm 1 and the optimal solution, respectively. If $S = Q$, then S is optimal and we have

[5] For more details, see http://www.slate.com/articles/technology/cover_story/2016/01/how_facebook_s_news_feed_algorithm_works.single.html.

[6] See https://www.addictivetips.com/web/get-continue-watching-on-top-in-netflix/ for details on where the "continue watching" content is displayed on Netflix.

proven the claim. Suppose instead that $S \neq Q$. If $Q \subset S$, then Q cannot be optimal, because S results in more searches. The case $S \subset Q$ is also not possible. Any $j \in Q \cap \overline{S}$ must have $L_j \geq L_k, \forall k \in S$, because Algorithm 1 always selects products with the lowest L_j. So any such j could always be appended to S at the end. Because Algorithm 1 stopped, it means there are no such products. Finally, if $|S| = |Q|$, even though $S \neq Q$ (given that Algorithm 1 does not produce a unique allocation), then the claim is also proven.

Then there must be at least one element in S that is not in Q and vice versa. Consider locations in increasing order of their search cost. Let l be the first location in which $\overrightarrow{j}^S$ differs from $\overrightarrow{j}^Q$, i.e. $j_l^S \neq j_l^Q$. It must be that $L_{j_l^S} \leq L_{j_l^Q}$, because Algorithm 1 chose to allocate j_l^S. Then for location l, the set of products not allocated as per Algorithm 1, $\overline{S}_l$ is equal to the set of products not allocated as per the optimal algorithm, $\overline{Q}_l$ except that the product j_l^S was swapped for j_l^Q, that is $\overline{S}_l = \overline{Q}_l \setminus \{j_l^S\} \cup \{j_l^Q\}$. Because $L_{j_l^S} \leq L_{j_l^Q}$, we can apply Lemma 1 to show that $C_l(\{L_j, j \in \overline{S}_l\}) \geq C_l(\{L_j, j \in \overline{Q}_l\})$. Therefore, by allocating j_l^S and discarding j_l^Q in the optimal solution, we obtain a solution $\overrightarrow{j}^{Q'}$ that results in $N_{Q'} \geq N_Q$ products being searched, and which differs from $\overrightarrow{j}^S$ by one less product. By repeatedly applying this transform, $\overrightarrow{j}^Q$ can be transformed to $\overrightarrow{j}^S$ with no decrease in total number of products searched. This shows that S is optimal. □

Proof of Theorem 4[7] Suppose Algorithm 2 produces a set of allocated products S, while the optimal algorithm produces a set Q. If $S = Q$, then S is optimal and we have proven the claim. It is also clear that neither $Q \subset S$ (contradicts Q being optimal), nor $S \subset Q$ (could append any $j \in Q \cap \overline{S}$ to S at the end) are possible. Then there must be at least one element in S that is not in Q and vice versa.

We first show that there exist feasible product allocations $\overrightarrow{j}^S$ and $\overrightarrow{j}^Q$, such that all products that are included in both S and Q are placed in the same location. Let j be placed in location l in $\overrightarrow{j}^S$ and in location l' in $\overrightarrow{j}^Q$. If $l < l'$, then replace j in l' in $\overrightarrow{j}^S$. Now, j is placed in the same location in the two allocations. If $l > l'$, the same transformation can be applied to $\overrightarrow{j}^Q$. Therefore, we can transform any two feasible product allocations into allocations for which all common products are placed in the same location.

Denote by ϕ_j the expected payoff of a product j. Let a be the highest expected payoff product that is included in S and not in Q, that is $a \in S, a \notin Q$. Then, it must be that $\phi_a \geq \phi_b, \forall b \in Q, b \notin S$. Otherwise, if $\phi_b > \phi_a$, Algorithm 2 would consider product b before product a, and include it in S.

Consider now the location at which product a is placed in $\overrightarrow{j}^S$. Let c be the product placed in the same location in $\overrightarrow{j}^Q$. Because all common elements are placed in the same location under both algorithms, we know that $a \in S, a \notin Q$ and $c \in Q, c \notin S$. Swapping c with a in Q cannot decrease the total payoff from Q because $\phi_a \geq \phi_c$,

[7] Our proof is adapted from Cormen et al. (2009)

and it cannot increase the total payoff from Q because Q is optimal. Thus, swapping c with a, gives a feasible product allocation that differs from the set S in one less product than did Q. Through repetition, the set Q can be transformed to S with no decrease in payoff. Therefore, the set S was optimal. □

Proof of Theorem 5[8] If $\overrightarrow{j}^S$ has allocated products with lower expected net utility to lower search cost locations, then the claim has been proven. Suppose instead that this does not hold. Then $\exists j, k \in S$, such that $\gamma_j < \gamma_k$, but k was located in a lower cost location than j. Without loss of generality, suppose k is located at l and j at $l + h$, where $h > 0$. Because $\gamma_j < \gamma_k$, then it follows that $L_j \leq L_k$. Product j is located at $l + h$, which means that $L_j \geq l + h$. Because $L_j \leq L_k$, then it must also hold that $L_k \geq l + h$. Thus, we can swap j and k in S. The resulting allocation is still feasible, because the consumer would be willing to search both products in their new locations. Also, this cannot increase the retailer's expected payoff, because S is optimal. Thus, if $\overrightarrow{j}^S$ has not allocated products with lower expected net utility to lower search cost locations, it can be arranged in such a way, while preserving optimality. □

References

Ainslie, A., & Rossi, P. (1998). Similarities in choice behavior across product categories. *Marketing Science*, *17*(2), 91–106.

Anglin, P. (1990). Disjoint search for the prices of two goods consumed jointly. *International Economic Review*, *31*(2), 383–408.

Anderson, S., & Renault, R. (1999). Pricing, product diversity, and search cost: a Bertrand-Chamberlin-Diamond model. *The RAND Journal of Economics*, *30*(4), 719–735.

Arbatskaya, M. (2007). Ordered search. *The RAND Journal of Economics*, *38*(1), 119–126.

Armstrong, M., Vickers, J., Zhou, J. (2009). Prominence and consumer search. *The RAND Journal of Economics*, *40*(2), 209–233.

Armstrong, M., & Zhou, J. (2011). Paying for prominence. *The Economic Journal*, *121*(556), 368–395.

Athey, S., & Ellison, G. (2011). Position auctions with consumer search. *The Quarterly Journal of Economics*, *126*(3), 1213–1270.

Branco, F., Sun, M., Villas-Boas, J.M. (2012). Optimal search for product information. *Management Science*, *58*(11), 2037–2056.

Branco, F., Sun, M., Villas-Boas, J.M. (2016). Too much information? Information provision and search costs. *Marketing Science*, *35*(4), 605–618.

Burdett, K., & Malueg, D. (1981). The theory of search for several goods. *Journal of Economic Theory*, *24*, 362–376.

Carlin, B., & Ederer, F. (2018). Search fatigue. *Review of Industrial Organization*. forthcoming.

Carlson, J., & McAfee, P. (1984). Joint search for several goods. *Journal of Economic Theory*, *32*, 337–345.

Chandon, P., Hutchinson, J.W., Bradlow, E.T., Young, S.H. (2009). Does in-store marketing work? Effects of the number and position of shelf facings on brand attention and evaluation at the point of purchase. *Journal of Marketing*, *73*(6), 1–17.

Chen, Y., & He, C. (2011). Paid placement: Advertising and search on the internet. *The Economic Journal*, *121*(556), 309–328.

Chen, Y., & Yao, S. (2016). Sequential search with refinement: Model and application with click-stream data. *Management Science*, *34*(4), 606–623.

[8]Our proof is adapted from Cormen et al. (2009).

Chib, S., Seetharaman, P., Strijnev, A. (2002). Analysis of multi-category purchase incidence decisions using IRI market basket data. *Econometric Models in Marketing*, *16*, 57–92.

Chick, S., & Frazier, P.I. (2012). Sequential sampling with economics of selection procedures. *Management Science*, *58*(3), 1–16.

Chintagunta, P., & Halder, S. (1998). Investigating purchase timing behavior in two related product categories. *Journal of Marketing Research*, *35*(1), 43–53.

Cormen, T., Leiserson, C., Rivest, R., Stein, C. (2009). *Introduction to algorithms*, 3rd edn. Cambridge: MIT Press. Chapter 16.5 "Task scheduling problem".

De los Santos, B., & Koulayev, S. (2017). Optimizing click-through in online rankings with endogenous search refinement. *Marketing Science*, *36*(4), 542–564.

Erdem, T., & Winer, R. (1999). Econometric modeling of competition: A multi-category choice-based mapping approach. *Journal of Econometrics*, *89*, 159–175.

Ellison, G., & Ellison, S. (2009). Search, obfuscation, and price elasticities on the internet. *Econometrica*, *77*(2), 427–452.

Gamp, T. (2017). Guided search. Working article.

Gatti, R. (1999). Multi-commodity consumer search. *Journal of Economic Theory*, *86*, 219–244.

Ghose, A., Ipeirotis, P., Li, B. (2012a). Surviving social media overload: Predicting consumer footprints on product search engines. Working article.

Ghose, A., Ipeirotis, P., Li, B. (2012b). Designing ranking systems for hotels on travel search engines by mining user-generated and crowd-sourced content. *Marketing Science*, *31*(3), 492–520.

Ghose, A., Ipeirotis, P., Li, B. (2014). Examining the impact of ranking on consumer behavior and search engine revenue. *Management Science*, *60*(7), 1632–1654.

Granbois, D. (1968). Improving the study of customer in-store behavior. *Journal of Marketing*, *32*(October), 28–33.

Haan, M., Moraga-Gonzalez, J., Petrikaite, V. (2018). A model of directed consumer search. *International Journal of Industrial Organization*, forthcoming.

Hagiu, A., & Jullien, B. (2011). Why do intermediaries divert search? *The RAND Journal of Economics*, *42*(1), 337–362.

Hansen, K., Singh, V., Chintagunta, P. (2006). Understanding store-brand purchase behavior across categories. *Marketing Science*, *25*(1), 75–90.

Hui, S., Inman, J., Huang, Y., Suher, J. (2013). The effect of in-store travel distance on unplanned spending: Applications to mobile promotion strategies. *Journal of Marketing*, *77*(March), 1–6.

Johnson, J. (2017). Unplanned purchases and retail competition. *American Economic Review*, *107*(3), 931–965.

Ke, T., Shen, Z.M., Villas-Boas, J.M. (2016). Search for information on multiple products. *Management Science*, *62*(12), 3576–3603.

Ke, T., & Villas-Boas, J.M. (2017). Optimal learning before choice. Working article.

Koulayev, S. (2014). Search for differentiated products: Identification and estimation. *The RAND Journal of Economics*, *45*(3), 553–575.

Larson, J., Bradlow, E., Fader, P. (2005). An exploratory look at supermarket shopping paths. *International Journal of Research in Marketing*, *22*(4), 395–414.

Levav, J., Heitmann, M., Herrmann, A., Iyengar, S. (2010). Order in product customization decisions: Evidence from field experiments. *Journal of Political Economy*, *118*(2), 274–299.

Liu, T. (2011). *Learning to rank for information retrieval*. Berlin: Springer.

Manchanda, P., Ansari, A., Gupta, S. (1999). The "shopping basket": A model for multicategory purchase incidence decisions. *Marketing Science*, *18*(2), 95–114.

Mehta, N. (2007). Investigating consumers' purchase incidence and brand choice decisions across multiple product categories: A theoretical and empirical analysis. *Marketing Science*, *26*(2), 196–217.

McAfee, P. (1995). Multiproduct equilibrium price dispersion. *Journal of Economic Theory*, *67*, 83–105.

Ngwe, D., Ferreira, K., Teixeira, T. (2019). The impact of increasing search frictions on online shopping behavior: Evidence from a field experiment. Working paper.

Petrikaite, V. (2018). Consumer obfuscation by a multiproduct firm. *The RAND Journal of Economics*, *49*(1), 206–223.

Redden, J.P. (2008). Reducing satiation: The role of categorization level. *Journal of Consumer Research*, *34*(5), 624–634.

Rhodes, A. (2011). Can Prominence matter even in an almost frictionless market? *The Economic Journal*, *121*(556), 297–308.

Seetharaman, P., Ainslie, A., Chintagunta, P. (1999). Investigating household state dependence effects across categories. *Journal of Marketing Research*, *36*(4), 488–500.

Seetharaman, P., Chib, S., Ainslie, A., Boatwright, P., Chan, T., Gupta, S., Mehta, N., Rao, V., Strijnev, A. (2005). Models of multi-category choice behavior. *Marketing Letters*, *16*(2), 239–254.

Shelegia, S. (2011). Multiproduct pricing in oligopoly. *International Journal of Industrial Organization*, *30*(2), 231–242.

Singh, V., Hansen, K., Gupta, S. (2005). Modeling preferences for common attributes in multicategory brand choice. *Journal of Marketing Research*, *42*(2), 195–209.

Song, I., & Chintagunta, P. (2006). Measuring cross-category price effects with aggregate store data. *Management Science*, *52*(10), 1594–1609.

Song, I., & Chintagunta, P. (2007). A discrete-continuous model for multicategory purchase behavior of households. *Journal of Marketing Research*, *44*(4), 595–612.

Stigler, G. (1961). The economics of information. *Journal of Political Economy*, *65*(3), 213–225.

Ursu, R. (2018). The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions. *Marketing Science*, *37*(4), 530–552.

Varian, H. (2007). Position auctions. *International Journal of Industrial Organization*, *25*(6), 1163–1178.

Weitzman, M. (1979). Optimal search for the best alternative. *Econometrica*, *47*(3), 641–654.

Wilson, C. (2010). Ordered search and equilibrium obfuscation. *International Journal of Industrial Organization*, *28*(5), 496–506.

Wolinsky, A. (1986). True monopolistic competition as a result of imperfect information. *The Quarterly Journal of Economics*, *101*(3), 493–512.

Yoganarasimhan, H. (2018). Search personalization using machine learning. *Management Science*, forthcoming.

Zhou, J. (2011). Ordered search in differentiated markets. *International Journal of Industrial Organization*, *29*(2), 253–262.

Zhou, J. (2014). Multiproduct search and the joint search effect. *American Economic Journal*, *104*(9), 2918–2939.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.