How Do Product Recommendations Help Consumers Search Products?

Evidence from a Field Experiment

Xiang (Shawn) Wan,[a] Anuj Kumar,[b] Xitong Li [c]

[a] Santa Clara University; [b] University of Florida; [c] HEC Paris

xwan@scu.edu; akumar1@ufl.edu; lix@hec.fr

Abstract

Product recommendations can benefit consumers' online product search via multiple underlying mechanisms, such as showing products that offer them high value, facilitating navigation on the website, or exposing more product information. However, it is unclear *ex-ante* which is the primary underlying mechanism that drives the benefits of product recommendations to consumers. We conducted a randomized field experiment to estimate the benefits of an item-based collaborative filtering (CF) recommendation system to consumers. We collect unique data on the affinity scores computed by an item-based CF algorithm to develop measures of a product's net value and horizontal (taste) fit for consumers. Our results indicate that product recommendations help consumers search for higher-value products that are lower-priced, fit their tastes better, or both. Besides that, we find that the ability to find higher-value products (rather than easy navigation or exposure to more product information) is the primary driver for consumers' higher purchase probabilities under recommendations. We further find a higher benefit of recommendations in product categories with higher price dispersion and heterogeneity in consumers' tastes, providing additional evidence for the lower price and better horizontal fit mechanisms. Finally, we find that when made available, consumers substitute their usage of other search tools on the website with product recommendations. Our findings have important implications for online retailers, policymakers, regulators, and item-based CF recommendation system design.

Keywords: Product recommendations; benefits of recommendations; consumer search; horizontal fit; field experiment

1. Introduction

The extant literature on product recommendations mainly examines their economic benefits to firms, such as their positive effect on product sales and how they can maximize product sales (De et al. 2010, Lee and Hosanagar 2019). Yet, the effect of product recommendations on sales and their benefits to consumers could come from multiple underlying mechanisms. Specifically, recommendation systems are designed to infer consumers' preferences based on their browsing behavior and accordingly recommend algorithmically identified related products. Thus, product recommendations are expected to help consumers find products that offer them higher net value, which could be due to (i) higher quality (vertical product attributes), (ii) better horizontal taste fit (match of horizontal product attributes with consumers' tastes), or (iii) lower price. However, recommendations also allow consumers to navigate directly from one product's page to another on the websites, inducing consumers to view more product pages and purchase more (*navigation effect*). Consumers view recommended products on a focal product's page and are thus exposed to more product information per product page view. Exposure to more product information could also drive higher sales under recommendations (*information exposure effect*).

The managerial implications of recommendation systems may differ based on which mechanism drives their effect. If product recommendations indeed help consumers find products of higher net value, the retailers could increase consumer satisfaction and retention by investing in and improving the recommender system algorithm. However, if the navigation effect is driving the benefits of product recommendations, retailers should offer more hyperlinked products on a product's page to allow greater opportunities for direct navigation across web pages. Finally, if the information exposure effect is the primary driver, retailers should optimize the product information on the web pages of their websites. To our knowledge, no prior study has unlocked these underlying mechanisms that could drive the benefits of product recommendations in a field study to provide clear managerial implications. We aim to fill this literature gap by examining whether product recommendations help consumers discover products of higher net value, while ruling out alternative mechanisms such as the pure navigation effect and information exposure effect.

While it is easy to observe the price of purchased products, their value and fit with consumers' tastes are latent (Ratchford and Srinivasan 1993). Due to this challenge, prior research has primarily examined the effect of search on the price of the purchased products (Ratchford and Srinivasan 1993, Seiler and Pinna 2017, Ursu et al. 2020), but not on the net value or horizontal fit of the purchased products to consumers. We utilize unique data on the "similarity" scores (affinity scores) computed between products by the recommendation algorithm to impute their net values to consumers. Using these estimates of product net value, we estimate the impact of recommendations in helping consumers search for and purchase products of higher net value. We further separately estimate whether the higher net value is attributable to lower product price, a better fit of the product's attributes with consumers' tastes, or both.

Besides product recommendations, e-commerce websites commonly offer two other search tools to consumers: keyword-based and product category-based searches. Prior literature on consumer search and product recommendations offers little empirical evidence on the substitution among the different search tools on e-commerce websites. Understanding the relative efficacies of these search tools could provide managerial insights on how to design them better and how much to invest in them.

In sum, we attempt to answer the following research questions: (i) Do product recommendations help consumers discover high-value products? (ii) How do product recommendations influence consumers' purchase probability and the price and horizontal fit of purchased products? (iii) How do product recommendations affect consumers' usage of other search tools on e-commerce websites?

Answering the above research questions poses both demand- and supply-side challenges. On the demand side, consumers endogenously choose their search efforts (product recommendation usage). Unobserved confounders (such as consumers' price sensitivity) may simultaneously drive consumers' product recommendation usage and search process outcomes, e.g., the price or horizontal fit of searched products. We address this challenge by exogenously manipulating the availability of product recommendations to consumers in a field experiment. In this experiment, we randomly assigned visitors to view one of the two versions of an online retailer's website. The website's treated version recommended four products that are most co-viewed and co-purchased with and in the product category of the focal product (hereafter FP) on the FP's page. These four products are called recommended products (hereafter RP).[1] The control version of the website hid the RPs on the FP's page. On the supply side, retailers may strategically recommend products to optimize profits or reduce inventory.[2] In such cases, the retailer-level unobserved factors may confound the estimate of the benefits of recommendations to visitors. Fortunately, the retailer in our research context did not strategically recommend products on its website.

The recommendation algorithm computes the weighted sum of co-views and co-purchases of the FP and its RP by past visitors (*affinity scores*). We collected data on the affinity scores of each FP-RP pair. Among various products viewed during the search process in a product category (*purchase funnel*), we find that visitors are more likely to buy products having higher affinity scores with the first FP in a purchase funnel.[3] Intuitively, visitors view/buy products that offer them high net values. Since visitors prefer to view high affinity score products, a product's affinity score with the first FP in our field setting offers a measure of its net value to the visitors. While a product's vertical quality is the same for all visitors, its horizontal fit with tastes may vary across visitors. Using this distinction, we estimate the part of a product's net value that captures its horizontal (taste) fit (see details about the measures of a product's net value and horizontal fit to visitors' taste in Section 5).

We measure a visitor's recommendation usage with (i) an indicator variable for the availability of recommendations and (ii) the total number of RP page views during a search process. Although the availability of recommendations to visitors is random, visitors endogenously choose their number of RP page views. We use the exogenous availability of recommendations to visitors as an instrumental variable to account for the endogeneity in their number of RP page views (search effort under recommendations).

[1] Recommendation algorithms based on such co-view (co-purchase) relationships between products belong to item-based CF recommendation algorithms. Details of the recommendation engine are presented in Section 3.2.

[2] For example, few recommendation systems are designed to maximize retailers' profits besides providing relevant recommendations to the users (Abdollahpouri et al. 2020, Zhang et al. 2021). Some recommendation systems balance these two objectives by recommending the high margin or high-value products from the pool of relevant products.

[3] Moreover, among the four RPs shown on the FP's page, we find that visitors are more likely to view and buy the RP that has a higher affinity score with the FP.

Then, we estimate the effect of recommendation usage by comparing the search and purchase behavior of treated and control visitors.

We find that treated visitors browse products of higher net value (as measured by the affinity score with the first FP in the purchase funnel), lower prices, and better horizontal fit than control visitors. The higher number of RP page views due to recommendations, on average, increases the net value and horizontal fit of browsed products. Moreover, additional RP page views under recommendations reduce the minimum, average, and median values, but not the maximum value, of the price distribution of searched products.

If recommendations help visitors find higher net value products, we should expect a higher purchase probability for treated visitors than control visitors. In line with this conjecture, we find that product recommendations increase purchase probability. We further find that the treated visitors purchase products with a higher likelihood because they could find products that are lower priced, have better horizontal taste fit, or both. For the converted visitors, we further estimate the effect of recommendations on purchased products. We find that the treated visitors purchase higher net value, lower-priced, and better horizontal fit products than control visitors. We conduct additional analyses in Section 6.2.5 to show that the benefit of recommendations is not due to the navigation or information exposure effects.

We leverage the variation in product characteristics across product categories to provide additional evidence for the lower price and better horizontal fit mechanisms. Our moderation analyses indicate that the effect of product recommendations is more salient in (i) finding lower-priced products for product categories with high average prices and relative price dispersions (ratio of variance over average price), where the scope of finding lower-priced products is larger; (ii) finding higher horizontal fit products for product categories with higher consumer taste heterogeneity (e.g., women's product categories), where the opportunity for finding better horizontal fit products is higher. These results further support that recommendations help consumers find lower-priced and better horizontal fit products.

Finally, we find evidence that visitors substitute other search tools with product recommendations, plausibly because recommendations help them search for higher net value products than other alternatives.

Economically, our estimates indicate that an additional RP page view under recommendations decreases 0.12 search page views and 0.65 product category page views, respectively.

Our paper makes several contributions to the literature on product recommendations. First, we provide empirical evidence that product recommendations help consumers search and buy products of higher net value that are lower priced, fit their tastes better, or both in a real-life business setting. Second, we find a higher benefit of recommendations in product categories with higher price dispersion and consumer taste heterogeneity, providing additional evidence for these underlying mechanisms. These results are new in the recommender systems literature. Lastly, our finding that consumers substitute the existing search tools with product recommendations is also new in the literature. The results highlight the relative efficacy of product recommendations compared to other online search tools.

2. Related Literature

Our research draws from two literature streams: the literature on product recommendations and the literature on consumer search costs and returns to search.

2.1. Related Literature on Product Recommendations

Our research is closely related to the literature on product recommendations. Many prior studies in this literature stream estimate the positive impact of product recommendations on product sales. De et al. (2010) conduct one of the earliest studies to show that using a recommender system could increase both promoted and non-promoted products' sales. Lee and Hosanagar (2019) find the positive impact of two types of collaborating filtering – purchase-based and view-based collaborative filtering (CF) – on sales volume, while the effect of purchase-based collaborative filtering is more pronounced. More recently, Lee and Hosanagar (2021) examine the heterogeneity in the positive effects of recommendations on purchase conversion rate and find larger effects on hedonic goods than utilitarian goods but not significantly different effects for experience and search goods.

Many product recommendations are developed based on co-view or co-purchase relationships between products (Goldenberg et al. 2012, Thorat et al. 2015). Emerging literature focuses on examining the effects of product recommendations from this perspective. Oestreicher-Singer and Sundararajan (2012) find that the visibility of a co-purchase relationship can lead to an average threefold amplification of the influence of complementary products on each other's demands. Kumar and Tan (2015) document that recommending products on the focal product pages can increase not only the sales of the focal products (direct effect) but also the sales of complementary products (spillover effect). While the jointly displayed products are largely complementary in the study of Kumar and Tan (2015) (e.g., apparel and accessories), Kumar and Hosanagar (2019) examine the situation where the focal product and the recommended products are substitutes. They find that a recommendation link increases the daily focal product page views by seven percent, reduces focal product sales conditional on the page views by 8.5 percent, and increases recommended products' sales by 24.5 percent. Lin et al. (2017) examine how the product recommendation network's diversity could influence the effect of recommendations on product sales.

A highly related work by Li et al. (2022) examines the causal paths through which the recommender systems lead to consumer purchases and find that the presence of personalized recommendations increases consumers' purchase probability by affecting the breadth and depth of consumers' consideration sets. In contrast, our study examines the underlying benefits to consumers that could explain recommender systems' positive impacts on consumers' purchase probability. We provide field experimental evidence that product recommendations help consumers find and buy higher net value products (with lower prices and/or better horizontal fit). Thus, our findings are quite different and novel, compared to other papers in this literature.

Prior studies in the product recommendation literature primarily focus on examining the economic benefits that product recommendations bring to firms and online retailers. These studies explore the effects of product recommendations on product sales and purchase conversion rate and how these effects differ with product characteristics, such as hedonic versus utilitarian goods (Kumar and Hosanagar 2019, Lee and Hosanagar 2021) or the category diversity of the product recommendation network (e.g., Lin et al. 2017). The only exception is the recent work that has examined the welfare properties of recommender systems for consumers (Zhang et al. 2021). Our study differs from this study as we propose a novel method of empirically estimating a product's net value and horizontal fit to consumers for item-based CF recommender systems. To our knowledge, no prior study has answered whether the benefits of product recommendations to consumers are primarily due to finding lower-priced products, better-fit products, or some alternative underlying mechanisms.

2.2. Related Literature on Search Cost and Return to Consumer Search

Our research is also related to the literature on consumers' search costs. Consumers' search for product information is nontrivial and often costly. One assumption could be that consumers perform a sequential search (Weitzman 1979, Reinganum 1982), i.e., consumers continue searching until the marginal cost exceeds the expected marginal benefit of an extra search. Most studies on consumer search in an online environment follow this assumption (Chapelle and Zhang 2009, Chen and Yao 2017, Ghose et al. 2019). Accordingly, in the equilibrium, consumers stop searching when the expected marginal benefit of search equals the marginal search cost, thereby allowing researchers to measure the return to consumer search.

It is challenging to estimate consumer search costs because it is not observed and needs to be inferred from other observed variables. Ratchford and Srinivasan (1993) conduct one of the earliest studies to estimate the monetary return to consumer search in terms of the consumer finding lower-priced products. Specifically, they found that the median consumer could save $17.76 by searching for a lower car price for an additional hour. In a later study, Honka (2014) estimates the range of search costs in the US auto insurance industry from $35 to $170. She finds the search cost to be a more important driver of consumer retention than switching costs and customer satisfaction. Accordingly, she proposes eliminating search costs as the main lever to increase consumer welfare in the auto insurance industry. Ngwe et al. (2019) find that under certain conditions, deliberately increasing search costs by varying website navigation elements, especially those associated with accessing discounted items, could increase online retailers' average selling prices and overall expected purchase probability.

Estimating consumers' search costs is also challenging because consumers' search efforts are endogenous. For example, more price-sensitive consumers exert more effort in searching for lower-priced products. Ursu et al. (2020) account for this endogeneity by explicitly modeling a consumer's decision on how much time she spends searching to estimate an average search cost of \$0.07 per minute. Ursu et al. (2020) further show that consumers are more likely to purchase from restaurants that they spend more time searching online. More closely related to our research, Seiler and Pinna (2017) estimate the causal return to consumers' search efforts using "path-tracking" data obtained from shopping carts equipped with radio-frequency identification (RFID) tags in a physical store environment. Like Ratchford and Srinivasan (1993), Seiler and Pinna (2017) estimate the monetary return to consumer search in terms of lower prices paid for purchased products. They find an additional minute of search results in a price saving of \$2.10.

The studies in this literature stream estimate the monetary return to consumer search in terms of the discovery of lower-priced products (Ratchford and Srinivasan 1993, Seiler and Pinna 2017). Our study contributes to this literature by estimating the return on consumer search under recommendations to discover the higher net value and horizontal fit products besides lower-priced products. Our study also examines how consumers choose among the available search tools on the websites. Our research provides causal empirical estimates of consumers substituting other search tools with product recommendations in a randomized field study. Besides, while Seiler and Pinna (2017) do not find heterogeneity in returns to consumer search across product categories, we find substantial heterogeneity depending on relative price dispersions across product categories. Moreover, to our knowledge, no prior research compares consumers' search behaviors when product recommendations are available versus when they are unavailable, except the laboratory study by Dellaert and Häubl (2012).

3. Field Setup

3.1. Website Description

We conducted a field experiment on a US mid-size online retailer's website. The retailer's annual sales are over \$400 million, and the online channel accounts for 10 percent of the total sales. The retailer offers over 35,000 products for sale on its website. The website organizes products into eight categories: women's clothing, men's clothing, etc. Each of these main categories is further subdivided into subcategories. For example, the main category of women's clothing consists of several subcategories, such as women's dresses, tops, shorts, and skirts.

The contents on the website are organized into four levels: home page, product category pages, product subcategory pages, and product description page (hereafter referred to as "product page"). Webpages on a higher level have hyperlinks to navigate visitors to the subsequent lower-level webpages. Figure 1 shows an example of a product page for a women's top (*focal product* or *FP*). The product page has a large FP image with detailed product information, such as regular and discounted retail prices, available colors and sizes, and product descriptions. For visitors presented with product recommendations, the FP's product page also displays four recommended products or RPs, as the smaller images of four women's tops under the heading "MORE OPTIONS" on the right in Figure 1. A visitor can reach the product page of an RP by clicking its image on the FP's page. After the visitor clicks to view the RP's product page, the RP becomes a new FP, and its product page will display another set of recommended products.

3.2. Collaborative Filtering Product Recommendations

The retailer uses IBM's Coremetrics digital recommendation engine to recommend the RPs on an FP's page. The recommendation algorithm selects the RPs based on two rules. First, the recommendation algorithm computes the affinity score between every two products on the retailer's website based on their co-views and co-purchases over the last 30 days. The affinity score consists of the following four component scores: (i) view-to-view score, which counts the number of times a visitor views both the FP and RP in the same session; (ii) view-to-buy score, which counts the number of times a visitor views the FP but buys the RP in the same session; (iii) buy-to-buy score, which counts the number of times a visitor purchases both the FP and RP but not necessarily in the same session; (iv) abandon-to-buy score, which counts the number of times a visitor abandons the carted FP to buy the RP in the same session. The algorithm of the IBM recommendation engine aggregates the four component scores using the following formula with defaulted weight: 70 (view-to-view score), 20 (view-to-buy score), 5 (buy-to-buy score), and 5 (abandon-to-buy score). The recommendation engine computes the affinity scores for each pair of products on the retailer’s website at 4 AM daily.

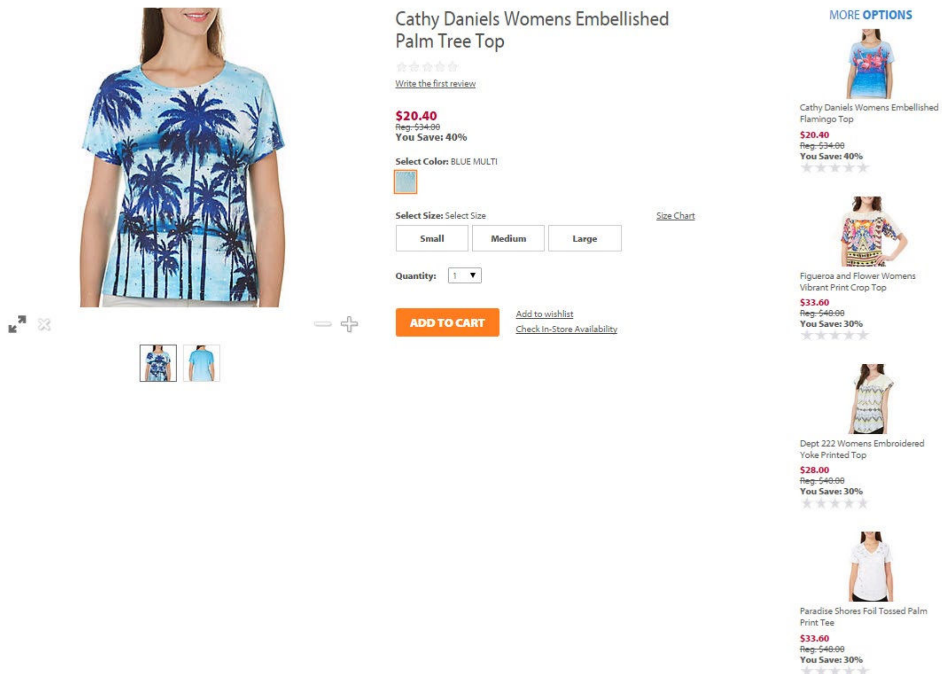

Figure 1. Product Page of an FP with four RPs

The above recommendation system utilizes the item-based collaborative filtering (CF) recommendation algorithm that infers the relationship between two products from the co-views and co-purchases of visitors on the retailer’s website (Sarwar et al. 2001, Linden et al. 2003). Item-based CF algorithms are among the most widely used recommendation algorithms in practice and studied in the literature (e.g., Adomavicius et al. 2019, Lee and Hosanagar 2021, Li et al. 2022). For example, “Amazon Personalize” provides similar item-based recommendations – “People who purchased this item also purchased…”

Besides the affinity score, the recommendation algorithm only selects RPs from the FP’s product subcategory. The algorithm stores the top 15 products with the highest affinity scores and the same product subcategory as the FP. Finally, the top four products with the FP’s highest affinity scores are recommended on the FP’s page.

3.3. Product Category Search and Keyword Search

As described in Section 3.1, the retailer’s website organizes products in different product categories and then under various subcategories within each product category. This hierarchical organization of products helps visitors conduct *product category-based searches* on the website. Besides the hierarchical organization of products, the retailer's website also offers a search bar on each webpage where visitors can enter their search keywords. The search algorithm identifies the products that match the search keywords and displays them on the search result page.

3.4. Experimental Setup

Visitors' product recommendations usage would be correlated to their unobserved characteristics. For example, a price-sensitive visitor would use recommendations more to search for lower-priced products. Therefore, we need an exogenous variation in the availability of product recommendations across visitors to estimate their effect on visitors' search outcomes. We accomplish this exogenous shock by conducting a randomized field experiment on the retailer's website. Specifically, the retailer created two versions of product pages on its website. The FP's page displayed four RPs identified by the recommendation system in the treated version of product pages, as illustrated in Figure 1. In contrast, the four RPs were suppressed (not displayed) on the FP's page in the control version of product pages. We ran the field experiment for nine weeks, from 8th April 2015 to 9th June 2015. In the experiment period, we randomly assigned half of the visitors on the website to the treated version (*treated visitors*) and the remaining half to the control version (*control visitors*). The recommendation engine can identify the visitors based on different information, such as cookies and IP addresses, and consistently assign visitors to the same experimental version in all visits.

3.5. Effect of Recommendations

Once a visitor reveals her preferences by selecting to view the product page of an FP on the website, the recommendation system algorithmically identifies and displays the related RPs on the FP's page. Thus, recommendations enable visitors to navigate directly from the FP's page to a related RP that is likely aligned with the visitors' preferences. This way, recommendations could help visitors view more products that offer them higher value from numerous alternatives available on the retailer's website.

In this research, we aim to explore whether product recommendations could help visitors find products that offer them higher net value. A product's net value to a visitor is equal to its value to them minus the price they pay for it. Prior literature (Sutton 1986, Kwark et al. 2017) suggests that the value of a product can be decomposed into two components: a value component from the vertical quality and the horizontal (taste) fit. Vertical product attributes (e.g., the quality of a product's raw material and craftsmanship) determine its vertical quality, for which preferences are the same across all consumers. In contrast, horizontal taste fit refers to the fit of product attributes (e.g., flavor and color) with individual consumers' tastes, which are heterogeneous and idiosyncratic across consumers.

4. Data Description and Summary Statistics

4.1. Purchase Funnels and Balance Checks

A visitor searches for desired products in a product category by exploring various relevant web pages, such as category/subcategory product pages, product detail pages, and search result pages on a website. If the visitor finds her wanted product, she will purchase it; otherwise, she ends her search. In line with the prior literature, we call a visitor's sequence of web page views during her product search in a product subcategory as a "purchase funnel" (Bleier and Eisenbeiss 2015).[4] A visitor's purchase funnel begins from a session when she starts searching for products in a product subcategory and ends when she purchases a product in that product subcategory or the end of our data period, whichever is earlier. After buying products in a subcategory, if a visitor again explores products in the same subcategory, it is considered a new purchase funnel. Thus, visitors can have multiple purchase funnels in a product subcategory in our data. A visitor exploring products in more than one product subcategory during a session will be in multiple purchase funnels.

We conduct analyses on purchase funnels under five main product categories. The product

[4] Visitors may go through multiple stages of the purchase process while purchasing a product in a product category. They search for product information, form a consideration set of products, evaluate products in the consideration set, and finally purchase the chosen product.

subcategories under these main categories are well-defined and only contain substitute products.[5] For example, women's skirts contain only women's skirts of all types. We dropped funnels in which visitors do not view a product page because we do not have information on those visitors' product searches. Ultimately, we have data of 573,665 purchase funnels across 69 product subcategories by 430,349 unique visitors. Among them, 287,374 (50.09 percent) funnels were for treated visitors (*treated purchase funnels*), and the remaining were for control visitors (*control purchase funnels*).

We check whether the random assignment of recommendations is valid in the final sample of visitors and purchase funnels. We conduct balance checks at both the visitor and purchase funnel levels. The results in Appendix A show statistically similar characteristics between the treated and control visitors (purchase funnels).

4.2. Page Views and Purchase Probability

Table 1 presents the summary statistics of purchase probability and different types of page views for the treated and control purchase funnels. While visitors purchase products in 6.1 percent of the treated purchase funnels, they do so in 6.0 percent of the control funnels. The difference in the two purchase probabilities is statistically significant, providing preliminary evidence that displaying product recommendations on FP pages can increase visitors' purchase likelihood.

In addition, we compute several variables related to visitors' search behaviors, such as the number of webpage views, FP page views, RP page views, category page views, and search result page views. We count each product's page view in a purchase funnel as an FP page view. If a product viewed in the purchase funnel appears as an RP (either explicitly in the treated version or hidden in the control version) on one of the earlier viewed products' pages, we count such page view as an RP page view.

[5] Since consumers search, evaluate, and choose a product among its substitute products in a purchase funnel, we dropped those product subcategories that also contain complementary or unrelated products in the same product subcategory. For example, under the product main category "Home", the product subcategory of "Home Décor" includes furniture, rugs and mats, sheets, and decorative pillows. It is not appropriate to consider the furniture, rugs, sheets, and pillows browsed by a visitor in her multiple sessions to be one purchase funnel. Accordingly, we dropped three main product categories.

Table 1. Summary Statistics of Page Views and Purchase Probability

	Control Purchase Funnels Mean (S.D.)	Treated Purchase Funnels Mean (S.D.)	Diff in Means (t-stats)
Purchase probability	0.060 (0.24)	0.061 (0.24)	0.001* (2.04)
# of webpage views	5.22 (8.74)	5.27 (8.67)	0.05* (2.17)
# of FP page views	1.68 (2.03)	1.83 (2.23)	0.14*** (25.36)
# of RP page views	0.21 (0.75)	0.36 (1.03)	0.16*** (66.28)
# of FP page views / # of total webpage views	0.65 (0.36)	0.66 (0.35)	0.01*** (8.97)
# of category page views	1.99 (5.83)	1.90 (5.65)	-0.09*** (-6.24)
# of search result page views	0.78 (2.63)	0.77 (2.61)	-0.02* (-2.53)

Notes: Diff. in means = Treated - Control. + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 1 reports the summary statistics, revealing several interesting comparisons between the two types of purchase funnels. First, we find a statistically higher average number of webpage views in treated versus control purchase funnels (5.27 versus 5.22). Second, we find a statistically higher average number of FP page views (1.83 versus 1.68) and RP page views (0.36 versus 0.21) in treated versus control purchase funnels. These results suggest that displaying related RPs on the FP's pages drives more FP and RP page views in treated purchase funnels plausibly because treated visitors can directly navigate to the related RP's page from the FP's page. Third and more interestingly, we find a statistically higher ratio of FP page views over the total number of webpage views in treated over control purchase funnels (0.66 versus 0.65). The higher ratio indicates that visitors with recommendations can view more product pages from the same total number of webpage views, suggesting a more efficient search under recommendations. Finally, we find a statistically lower average number of category page views (1.90 versus 1.99) and search result page views (0.77 versus 0.78) in treated versus control purchase funnels. These results indicate that treated visitors substitute their usage of alternative navigation tools (product category-based and keyword-based search) due to the recommendations on FP's pages.

4.3. Model-Free Evidence for Lower Price

Table 2 reports the summary statistics of the price distribution of products browsed in treated and control purchase funnels. We find a significantly lower mean value of average prices for products browsed in treated purchase funnels than control purchase funnels ($39.16 versus $39.35). We further find significantly lower mean values of the minimum prices, the prices at the 25th percentile, and the median prices of browsed products in treated purchase funnels. However, we do not find a significant difference in the mean values of maximum prices or the prices at the 75th percentile of browsed products between the two types of purchase funnels. These results provide model-free evidence that while treated visitors could find lower-priced products during the product search than control visitors, they view similarly high prices with and without recommendations.

Table 2. Summary Statistics of Price Distribution of Products Searched

	Control Purchase Funnels Mean (S.D.)	Treated Purchase Funnels Mean (S.D.)	Diff in Means (t-stat)
Average	39.35 (24.52)	39.16 (24.64)	-0.19^{**} (-2.89)
Minimum	37.93 (24.55)	37.56 (24.64)	-0.37^{***} (-5.72)
25^{th} percentile	38.21 (24.49)	37.89 (24.59)	-0.32^{***} (-4.91)
Median	39.26 (24.55)	39.06 (24.67)	-0.20^{**} (-3.05)
75^{th} percentile	40.46 (25.06)	40.40 (25.27)	-0.06 (-0.87)
Maximum	41.01 (25.45)	41.03 (25.73)	0.02 (0.31)

Notes: Diff. in means = Treated - Control. + p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001.

5. Measures of Net Value and Horizontal Fit

5.1. Net Value Measure

We collected unique data on daily affinity scores between each pair of products computed by the recommendation algorithm on the retailer's website. Below, we explain how we develop the measure of a product's net value to an average visitor based on affinity scores.

The affinity score between two products is the weighted sum of their number of co-views and co-purchases by past visitors to the website. For a visitor who views product A's page, if the affinity score between Products B and A is higher than between Products C and A, then the visitor is more likely to choose to view B than C by the definition of affinity scores. But visitors will view products that offer them high net value (utility), all else equal. Therefore, if visitors who view A prefer to view B which has a higher affinity score (with A) than C, B should offer, on average, a higher net value to the visitors. In other words, the affinity score of B (with A) correlates to its net value to visitors who view A. Thus, the affinity score of a product (with A) measures its net value to visitors who view A. Below, we conduct two empirical tests to show the validity of the measure.

Our first empirical test shows that an RP with a higher affinity score with its FP is more likely to be viewed (purchased) in the purchase funnel. While four RPs are explicitly displayed on an FP's page in the purchase funnel for treated visitors, these RPs are hidden on the FP's page to control visitors. We estimate the likelihood of each RP's view (purchase) on an FP's page in the purchase funnel with the following specification:

$$Y_{frit} = \alpha_f + \alpha_t + \beta \times LogAffScore_{frt} + \delta \times X_{frt} + \varepsilon_{frit}; \quad (1)$$

where *f*, *r*, *i*, and *t* denote the FP, RP, purchase funnel, and days. The dependent variables (Y_{rfit}) are the indicator variables for RP view and purchase. $LogAffscore_{frt}$ denotes the log of affinity score between the FP-RP pair on day *t*.[6] The control variable (X_{frt}) denotes RP's position on FP's page (Position=1,2,3, or 4).[7] We also include the FP- and day-fixed effects in the specification.[8] First, we estimate the specification (1) separately for the treated and control funnels. Then, we estimate it on the pooled data for both funnels with the indicator variable for treated funnels as an additional right-hand side variable. Table 3 reports the estimation results.

Columns (1)-(3) in Table 3 report estimates for the likelihood of an RP view. We first note that the signs of the estimated coefficients of the RP's position and whether the RP is explicitly displayed are consistent with our expectations and thus provide face validity to the empirical results.[9] Specifically, a lower RP's placement on the FP's page (higher RP's position) is associated with a lower likelihood of the

[6] We use the log transformation of affinity scores to account for its highly skewed distribution.

[7] If a control visitor views an FP's page, four RPs could still be selected based on their high affinity scores with the FP, but the RPs are just hidden on the FP's page. Accordingly, the RP's position in the control group can be defined in a similar way as in the treated group.

[8] In a robustness check, we further control the RP fixed effects and find qualitatively similar results.

[9] The correlation between an RP's position on the FP's page and its affinity score with the FP is -0.19, indicating a weak correlation between these two variables. Following prior literature (Burtch et al. 2013, Singh et al. 2014), we test the multicollinearity with the variance inflation factor (VIF) approach. The VIFs for affinity score and RP position in all regressions are much lower than the cutoff of 10 (Hadi and Chatterjee 2015) (Page 250). These analyses indicate that multicollinearity is not a concern in our analyses.

RP view, and the explicit display of the RP on the FP's page is associated with a higher probability of the RP view. Notably, we find positive and significant coefficients for affinity scores in all specifications, indicating that the probability of the RP view increases with its affinity score with the FP. We find qualitatively similar results for the likelihood of the RP purchase in columns (4)-(6) in Table 3. The results in Table 3 provide consistent evidence that an RP of a higher affinity score with the FP is more likely to be viewed and purchased.

Table 3. The Effect of RP's Affinity Score on RP View and Purchase

	DV: Whether RP View			DV: Whether RP Purchase		
	(1) Treated	(2) Control	(3) Pooled	(4) Treated	(5) Control	(6) Pooled
Affinity Score	0.0023***	0.0012***	0.0018***	0.0002***	0.0001*	0.0001***
	(0.0001)	(0.0001)	(0.0001)	(0.0000)	(0.0000)	(0.0000)
RP's Position	-0.0158***	-0.0103***	-0.0131***	-0.0014***	-0.0012***	-0.0013***
	(0.0002)	(0.0001)	(0.0001)	(0.0000)	(0.0000)	(0.0000)
Rec. Indicator			0.0271***			0.0009***
			(0.0002)			(0.0001)
FP fixed effect	Yes	Yes	Yes	Yes	Yes	Yes
Day fixed effect	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.0390	0.0394	0.0368	0.0124	0.0142	0.0088
Obs.	2,071,179	1,899,339	3,970,520	2,071,179	1,899,339	3,970,520

Notes: + p < .10, * p < .05, ** p < .01, *** p < .001

A visitor selects to view the product page of the first FP from among many products she views on the other web pages (e.g., the home page, product category main page, etc.). Therefore, the first FP viewed in a purchase funnel reveals the visitor's general preference in a product category. Subsequent products (FPs) in the purchase funnel having high-affinity scores with the first FP should also be aligned to visitors' preferences and offer them high net values. Accordingly, our second empirical tests show that products having higher affinity scores with the first FP in a purchase funnel are more likely to be purchased. Accordingly, we compute the affinity scores for all viewed products with the first FP in a purchase funnel. Visitors select to view the first FP among other products they viewed in its product category. Therefore, we compute the first FP's affinity score as the average of its affinity scores with other products in its product subcategory.[10] We estimate the effect of products' affinity scores on their purchase probability with the following specification:

$$PurProb_{fit} = \alpha_i + \alpha_t + \beta \times LogAffScore_{fit} + \delta \times X_{fit} + \varepsilon_{fit}; \qquad (2)$$

where f, i, and t, respectively, denote products (FPs), purchase funnels, and days. The dependent variable ($PurProb_{fit}$) is the probability of an FP purchase in funnel i on day t. $LogAffscore_{fit}$ denotes the log of affinity score of product f in funnel i on day t. The control variable (X_i) denotes product f's order (or position) in the purchase funnel. We also include product funnel and day-fixed effects in the specification (2).[11] First, we estimate the specification (2) separately for the treated and control funnels. Then, we estimate it by pooling the data for both funnels with the indicator variable for treated funnels as an additional right-hand side variable. Table 4 reports the estimation results.

Table 4. The Effect of FP's Funnel-level Affinity Score on FP Purchase

	DV: Whether FP Purchase		
	(1) Treated	(2) Control	(3) Pooled
Affinity Score with the first FP	0.0011***	0.0010***	0.0010***
	(0.0001)	(0.0001)	(0.0001)
FP's Position in the Funnel	0.0044***	0.0046***	0.0045***
	(0.0001)	(0.0001)	(0.0001)
Rec. Indicator			*Dropped*
Funnel fixed effect	Yes	Yes	Yes
Day fixed effect	Yes	Yes	Yes
R^2	0.4542	0.4754	0.4645
Obs.	482,135	524,990	1,007,125

Notes: + p < .10, * p < .05, ** p < .01, *** p < .001

We first note that the signs of the estimated coefficients of the FP's position (order) are significantly positive, suggesting that products viewed later in a purchase funnel are more likely to be purchased. This result is consistent with the prior findings that consumers are more likely to buy products viewed in the later stages of a purchase funnel (Lambrecht and Tucker 2013, Bleier and Eisenbeiss 2015), thus providing

[10] A product that has a higher average affinity with other products in its product subcategory has high visibility and thus, higher likelihood of being viewed first. We find qualitatively similar results from specification (2) with several alternative affinity scores for the first FP, such as not considering the affinity score of the first FP or keeping it as zero.

[11] In a robustness check, we include additional control variables in the estimation, including the average price of the RPs and the average affinity score of the RPs with the FP. We find qualitatively similar results.

face validity to the empirical results. The coefficients of affinity scores are all positive and statistically significant, indicating that FPs that have a higher affinity score with the first FP in a purchase funnel are more likely to be purchased.

The results in Tables 3 and 4, together, provide consistent evidence that visitors are more likely to view and purchase higher affinity score products. Since visitors would view and purchase products that offer them high net value, these findings indicate that the affinity score of an FP (with the first FP) viewed in a purchase funnel measures the FP's net value to visitors.

5.2. Horizontal Fit Measure

In this section, we employ the residual approach to measure an FP's horizontal (taste) fit, a component of its total net value to the visitors. The labor economics and financial accounting literature have widely used the residual approach to decompose a theoretical construct into components unrelated to certain factors (Solow 1956, 1957, Jones 1991, Jorgenson and Stiroh 1999, Basu et al. 2006, Dou et al. 2013, Kim et al. 2014, Ali and Zhang 2015, Brynjolfsson et al. 2021). Specifically, as described in Section 3.5, a product's value to a consumer can be decomposed into two components: its vertical quality and horizontal fit (Sutton 1986, Kwark et al. 2017). Following the recent work by Shi and Raghu (2020), we write the net value (utility) of a focal product *f* to consumer *k* as the specification (3.1):

$$NetValue_{kf} = \text{Vertical quality } (V_f) + \text{Horizontal fit of } f\text{'s design and consumer } k\text{'s idiosyncratic taste } (H_{kf}) - \text{Price } (\mathrm{P}_f)\ ; \tag{3.1}$$

Thus, the net value of a product *f* that consumer *k* views in purchase funnel *i* on day *t* can be written as:

$$NetValue_{kift} = \mathrm{V}_f + \mathrm{H}_{kf} - \mathrm{P}_{ft}\ ; \tag{3.2}$$

Note that vertical quality is uniform across consumers, whereas horizontal fit of product *f* varies across consumers. Since the retailer may offer price promotions on certain days, we allow the product prices to vary across days in the specification (3.2). As explained in the previous section, we use the affinity score of product *f* as the measure of its net value. Since the vertical quality of a product is constant across consumers, it could be identified with the product-fixed effects in the econometric specification. Accordingly, we estimate the following specification (3.3):

$$AffinityScore_{kift} = \beta_f + \beta_1 \times Price_{ift} + \varepsilon_{kift}; \quad (3.3)$$

where $AffinityScore_{kift}$ indicates the affinity score of product *f* viewed by consumer *k* in purchase funnel *i* on day *t*. $Price_{ift}$ indicates the price of product *f* in purchase funnel *i* on day *t*. β_f indicates product fixed effects, which accounts for the vertical quality of product *f*.[12] β_1 captures the average price sensitivity of consumers. ε_{kift} is the residual which captures the variation in $AffinityScore_{kift}$ after controlling for vertical quality and price effect of the product. Since a product's affinity score measures its net value, the residual ε_{kift} measures the part of its net value that varies across consumers and is uncorrelated to its vertical quality and price. Thus, the residual component of the product *f*'s net value captures the extent to which product *f* fits consumer *k*'s idiosyncratic taste, a measure of horizontal taste fit. Note that our measures of a product's net value and horizontal fit are with respect to the first FP in the purchase funnel because they are developed based on the product's affinity score with the first FP.

Since our estimation of the effects of product recommendations is on the purchase funnel level, we compute the average values of $AffinityScore_{if}$ and $HFit_{if}$ across all viewed products in purchase funnel i, denoted by $AvgAffScore_i$ and $AveHFit_i$, respectively. Table 5 reports their summary statistics for treated and control purchase funnels. These results show that the products browsed in treated purchase funnels on average have a higher affinity score and a higher value of horizontal fit than control purchase funnels. Thus, Table 5 results provide preliminary evidence that product recommendations, on average, help visitors find and view products with higher net value and better horizontal fit with their tastes.

Table 5. Summary Statistics of *AveAffScore* and *AveHFit*

	Control Purchase Funnels Mean (S.D.)	Treated Purchase Funnels Mean (S.D.)	Diff in Means (t-stat)
Avg. Affinity Score	1.285 (2.316)	1.678 (2.720)	0.393^{***} (59.00)
Avg. Horizontal Fit	0.187 (1.515)	0.452 (1.792)	0.266^{***} (60.54)

Notes: Diff. in means = Treated - Control. + p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001.

[12] Note that $Price_{ift}$ in the specification (3.3) can be further decomposed into two parts: original price that is constant across consumers over time, and the discounts that may vary over time. Accordingly, β_f in the specification (3.3) can capture the effect of vertical quality and the time-invariant original price.

6. Results

6.1. Effect of Recommendations on Consumers' Search Behavior

We examine the effect of recommendations on visitors' search behavior on the website with the following specifications:

$$Y_i = \beta_c + \beta_1 \times Rec_i + \varepsilon_i; \quad (4.1)$$

$$Y_i = \beta_c + \beta_1 \times NRPView_i + \varepsilon_i; \quad (4.2)$$

where i denotes purchase funnels, and c denotes product categories. Y_i denotes visitors' search behavior in terms of (i) average affinity score, average price, and average horizontal fit of the products browsed in a purchase funnel and (ii) different price-related variables, including minimum price; prices at 25th, 50th, and 75th percentiles; and maximum price of price distribution in purchase funnel i. Rec_i denotes the indicator variable for the treatment of recommendations. β_c denotes the product category fixed effects that account for differences in the impact of unobserved product category-level factors on the dependent variable. Thus, β_1 estimates the treatment – *intent to treat* – effect of recommendations.

As discussed in Section 3.5, recommendations help visitors find products of higher net value by enabling them to view more relevant RP page views. Thus, we estimate the effect of the number of RP page views in purchase funnel i ($NRPView_i$) on visitors' search behavior in the specification (4.2). Although recommendations are randomly assigned, visitors endogenously choose their recommendation usage in purchase funnels. Therefore, $NRPView_i$ is endogenous in the specification (4.2). We account for endogeneity by instrumenting the number of RP page views with the recommendation indicator (Rec_i) and estimating the specification (4.2) using a two-stage least square regression (2SLS). Variable Rec_i satisfies the exogeneity condition for instrument variable (IV) due to the random assignment of recommendations across visitors. Moreover, Rec_i should significantly affect the number of RP page views to satisfy the relevance condition. As expected, we find an *F-value* of 142.8 for the exclusion of Rec_i in the first stage regression of $NRPView_i$ on Rec_i. The estimated *F-value* is significantly higher than the threshold value of 10 for the weak instrument, indicating that Rec_i satisfies the relevance condition for the IV (Bound et al. 1995, Dinkelman 2011). The second stage regression estimates the effect of the number of RP page views due to recommendations on visitors' price search behavior. Table 6 reports the estimated coefficients from specifications (4.1) and (4.2).

Table 6. Recommendations and Consumers' Search Behavior
Panel A. Consumer Search

	(1) Avg. Affinity Score	(2) Avg. Price	(3) Avg. Horizontal Fit
	Fixed Effect Specification (4.1)		
Rec. Indicator	0.388^{***} (0.018)	-0.122^{*} (0.048)	0.266^{***} (0.013)
R^2	0.119	0.457	0.009
	Fixed Effect 2SLS Specification (4.2)		
No. of RP Page View	2.470^{***} (0.154)	-0.773^{**} (0.282)	1.689^{***} (0.093)
Obs.	573,665	573,665	573,665

Notes: + p < .10, * p < .05, ** p < .01, *** p < .001. Standard errors cluster corrected at product subcategory level in parentheses.

Panel B. Price Search

	(1) Minimum Price	(2) Price at 25th Percentile	(3) Median Price	(4) Price at 75th Percentile	(5) Maximum Price
	Fixed Effect Specification (4.1)				
Rec. Indicator	-0.305^{***} (0.051)	-0.252^{***} (0.047)	-0.131^{**} (0.047)	0.009 (0.054)	0.085 (0.054)
R^2	0.451	0.453	0.455	0.442	0.431
	Fixed Effect 2SLS Specification (4.2)				
No. of RP Page View	-1.937^{***} (0.310)	-1.599^{***} (0.288)	-0.834^{**} (0.277)	0.056 (0.345)	0.540 (0.363)
Obs.	573,665	573,665	573,665	573,665	573,665

Notes: + p < .10, * p < .05, ** p < .01, *** p < .001. Standard errors cluster corrected at product subcategory level in parentheses.

We first estimate the recommendations' impact on visitors' search behavior on the website using the specification (4.1). We capture visitors' search behavior in terms of average price, average affinity score, and average horizontal fit of the products browsed in a purchase funnel. Panel A of Table 6 reports the effects of our treatment variable (*Rec*. Indicator) on the different outcomes. Column (1) shows that recommendations help visitors find higher affinity score products that offer them higher net value. The results in columns (2) and (3) suggest that recommendations help visitors find lower-priced products and better horizontal fit products. Hence, the higher average net values of products browsed under recommendations are due to their lower average prices, a better fit of their attributes with visitors' tastes, or both.

Panel B of Table 6 reports the effect of recommendations on the price distribution of the products browsed in a purchase funnel from the specification (4.2). Interestingly, recommendations have a negative effect on the minimum price, the price at the 25 percentiles, and the median price. However, the effect of recommendations on the price at the 75 percentile and the maximum price is insignificant. These results reveal that recommendations help visitors explore lower-priced products.

6.2. Benefits of Recommendations to Consumers

Analyses in the previous section reveal that recommendations help visitors find products that offer them higher net value, perhaps due to lower prices, a better fit with their tastes, or both. We expect that discovering higher-value products should affect visitors' purchase behaviors with recommendations. We first examine the effects of recommendations on visitors' purchase probability and then on the outcomes related to the purchased products.

6.2.1. Effect of Recommendations on Purchase Probability

We estimate the following specifications to evaluate the effect of recommendations on the probability of product purchase:

$$ProbPurchase_i = \beta_c + \beta_1 \times Rec_i + \varepsilon_i; \tag{5.1}$$

$$ProbPurchase_i = \beta_c + \beta_1 \times X_i + \varepsilon_i; \tag{5.2}$$

where $ProbPurchase_i$ denotes whether there is a product purchase in purchase funnel i or not. X_i denotes average affinity score, average horizontal fit, and the variables capturing price distribution of products browsed in purchase funnels, including the minimum price, the price at the 25^{th} percentile, and the median price of products browsed in the purchase funnel i. All other variables have the same meanings as in the previous specifications.

We first use the OLS specification (5.1) to estimate the effect of the intent to treat (*Rec. Indicator*) on visitors' purchase probability; the Logistic regression produces qualitatively similar results. Column (1) of Table 7 indicates that showing product recommendations on average results in a higher purchase probability. The increase in the visitors' purchase likelihood may be because they can find higher net value products with the help of recommendations.

To empirically show this fact, we examine the effect of average affinity score, average horizontal fit, and the variables capturing price distribution of products browsed in purchase funnels on the purchase probability using the specification (5.2). Since these variables related to visitors' product search are endogenous, we use the randomized treatment indicator variable as an instrumental variable and use the two-stage least square (2SLS) regression to estimate their effects. Since *Rec. indicator* is our randomized treatment variable, it satisfies the exclusion restrictions for the average affinity score in the specification (5.2). Column (2) in Table 7 reports the results from the 2SLS regressions. We find a positive and significant coefficient for the average affinity score, indicating that higher affinity scores of browsed products under recommendations indeed lead to a higher purchase probability.

Table 7. Effects of Recommendations on Purchase Probability

	(1) OLS	(2) 2SLS	(3) 2SLS	(4) 2SLS	(5) 2SLS	(6) 2SLS	(7) 2SLS
Rec. Indicator	0.001^{*} (0.001)						
Avg. Affinity Score		0.003^{*} (0.001)					
Avg. Horizontal Fit			0.005^{*} (0.002)				
Avg. Price				-0.010 (0.006)			
Minimum Price					-0.004^{*} (0.002)		
Price at 25th Percentile						-0.005^{*} (0.002)	
Median Price							-0.009^{+} (0.005)
R^2	0.0064						
Obs.	573,665	573,665	573,665	573,665	573,665	573,665	573,665

Notes: + p < .10, * p < .05, ** p < .01, *** p < .001. Standard errors cluster corrected at product subcategory level in parentheses.

We further estimate the effect of the discovery of better-fit and lower-priced products due to recommendations on purchase probability from the specification (5.2). We acknowledge that the *Rec. indicator* may not satisfy the exclusion restriction for horizontal fit and price variables individually in the specification (5.2). We present these estimates merely as suggestive evidence. We find a positive and significant estimate for $HFit_i$ in column (3) of Table 7, indicating that a higher taste fit under recommendations is associated with increased purchase probability. The results in columns (4)-(7) show that while the change in the average price of the browsed products does not have a significant effect on purchase probability, a decrease in the minimum price, the price at the 25 percentiles, and the median price of the browsed products results in a higher purchase probability.

Overall, the results suggest that the positive effect of recommendations on purchase probability appears to be driven by the fact that visitors can find and view products of higher net value that are lower priced, have a better fit, or both.

6.2.2. Effect of Recommendations on Purchased Products

We estimate the effect of recommendations on the purchased products in this section using the specification (6). Specifically, we estimate the impact of recommendations on the affinity score, price, and horizontal fit of the purchased products:

$$Y_i = \beta_c + \beta_1 \times Rec_i + \varepsilon_i ; \quad (6)$$

where the dependent variable Y_i includes the affinity score (*PurAffScore*$_i$), price (*PurPrice*$_i$), and horizontal fit (*PurHFit*$_i$) of the purchased products. All other variables have the same meanings as in previous specifications. Table 8 reports the results from the specification (6). We find that recommendations help visitors purchase products of higher net value (as measured by the affinity score), lower price, and better fit to their tastes.

Table 8. Effects of Recommendations on Purchased Product

	(1) Purchase Affinity Score	(2) Purchase Price	(3) Purchase Horizontal Fit
Rec. Indicator	0.571^{***} (0.050)	-0.400^{***} (0.104)	0.487^{***} (0.033)
R^2	0.037	0.453	0.013
Obs.	34,859	34,859	34,859

Notes: + p < .10, * p < .05, ** p < .01, *** p < .001. Standard errors cluster corrected at product subcategory level in parentheses.

The results in Tables 7 and 8 show that visitors benefit from recommendations in two ways. First, they have a lower likelihood of failed search efforts due to a higher probability of product purchase under recommendations. Second, visitors with recommendations can search and purchase products that offer them higher net value due to lower product prices, better fit with their tastes, or both.

6.2.3. Robustness Checks

Since visitors view and buy more products under recommendations, affinity scores computed based on visitors' browsing behavior under recommendations may be inflated. One concern is that the effect of recommendations in discovering higher net value products based on inflated affinity scores may be upward biased. However, the random assignment of recommendations across visitors makes our findings robust to this possibility, as we explain below.

We identify the effect of recommendations in finding higher net value products by subtracting the average affinity scores of products viewed in treated purchase funnels from that in control funnels. The affinity score between a pair of products, computed by the recommendation algorithm, is not visible to consumers. More importantly, the random assignment of recommendations is at the visitor level, independent of the computation of the affinity scores between products. Any possible "inflation" in the affinity scores should not affect the difference in the average affinity scores of the products viewed/purchased in the treated and control purchase funnels. Therefore, our finding should be unbiased regardless of whether affinity scores were computed based on visitors' data with or without recommendations.

To further show that our results are robust regardless of whether we compute affinity scores based on settings with or without recommendations, we conduct the following robustness checks and report the detailed results in Appendix C. First, we estimate our results based on data from days 31-63 of the experiment and find qualitatively similar results in Appendix C1. The affinity scores in this duration were computed based on the equal proportion of visitors' browsing with and without recommendations. If inflation of affinity scores based on browsing under recommendations was biasing our result, we should have obtained different results in this analysis. Second, we recompute the affinity scores between products based on the browsing behavior of only control visitors (i.e., those who do not see recommendations) during our experiment period. We re-estimate the effect of recommendations on the net value of viewed/purchased products from recomputed affinity scores and find qualitatively similar results in Appendix C2.

6.2.4. Additional Evidence for Lower Price and Better Horizontal Fit with Recommendations

Next, we utilize the variation in product characteristics across product categories to provide additional evidence of better horizontal fit and lower prices of purchased products under recommendations. If recommendations help visitors find lower-priced products in a product category, this effect should be more prominent in the product categories with higher average prices and higher price dispersions. Similarly, if recommendations help visitors find better horizontal fit products, this effect should be more pronounced in product categories with highly heterogeneous visitors' tastes, such as women's apparel. We examine the moderating effect of product categories with the following specification (7):

$$Y_i = \beta_c + \beta_1 \times Rec_i + \beta_2 \times (Rec_i \times PCat_c) + \varepsilon_i ; \quad (7)$$

where the dependent variable Y_i is the price of purchased products. $PCat_c$ denotes the product category-level moderating factors in product subcategory *c*. We use the average product price and relative price dispersion (measured by the ratio of price variance to average price) in a product subcategory to show the differential effects of recommendations in finding lower-priced products. All other variables have the same meanings as in the previous specifications.

Table 9. Moderating Effects of Category-level Average Price and Price Dispersion

DV: PurPrice	(1)	(2)
Rec. Indicator	0.340	0.111
	(0.218)	(0.153)
Rec. Indicator × Average Price	-0.020***	
	(0.005)	
Rec. Indicator × Relative Price Dispersion		-0.072***
		(0.014)
R^2	0.453	0.453
Obs.	34,859	34,859

Notes: + p < .10, * p < .05, ** p < .01, *** p < .001. Standard errors cluster corrected at product subcategory level in parentheses.

Columns (1) and (2) in Table 9 report the estimated coefficients. We find negative and significant coefficients for the interaction term of *Rec. Indicator* with moderating factors related to price distribution. First, we find a significant negative coefficient for the interaction of *Rec. Indicator* and average product price, suggesting that visitors pay a lower purchase price by using recommendations for product categories with a higher average price. It is understandable because product categories with higher prices have more scope for finding products with lower prices. Second, we find a significant negative coefficient for the interaction of *Rec. Indicator* and relative price dispersion, suggesting that visitors pay a lower purchase price by using recommendations in product categories with a higher relative price dispersion. It is perhaps because such product subcategories offer visitors a greater opportunity to search for lower-priced products due to higher price dispersion across products.

Next, we examine the moderating effect of women's product category on the impact of recommendations in finding products that fit visitors' tastes better. First, we separately estimate the specification (7) on products in women's and men's apparel categories. The dependent variable in this analysis is the horizontal fit of purchased products. Next, we estimate the specification (7) on the pooled data for both product categories, including an indicator variable for women's apparel as $PCat_c$ variable in the specification (7).

Table 10. Moderating Effects of Category-level Taste Heterogeneity

DV: HFit	(1) Women	(2) Men	(3) Pooled
Rec. Indicator	0.535***	0.357***	0.357***
	(0.039)	(0.096)	(0.091)
Rec. Indicator × Women's Category			0.178^{+}
			(0.099)
R^2	0.008	0.009	0.008
Obs.	14,322	3,985	18,307

Notes: + p < .10, * p < .05, ** p < .01, *** p < .001. Standard errors cluster corrected at product subcategory level in parentheses.

Columns (1) and (2) of Table 10 report the results on the horizontal fit of purchased products for women's and men's apparel categories separately. We find positive and significant coefficients of *Rec. Indicator* for both product categories, while the estimated effect for women's apparel categories is larger. Moreover, column (3) of Table 10 reports a positive and significant coefficient for the interaction term, suggesting a higher increase in the horizontal fit of purchased products for women's apparel categories than for men's under recommendations. Perhaps because of the higher heterogeneity in visitors' tastes for products in women's compared to men's product categories, there is greater scope for recommendations to help visitors find products that better fit their tastes in the women's product categories.

6.2.5. Alternative Explanations

In this section, we rule out the possibility that the positive effects of recommender systems in our context are due to the pure navigation or information effects of recommendations.

6.2.5.1 Pure Navigation Effect

In our research context, by recommending related product on the FP's pages, the recommendation system provides visitors with direct navigation to another product's page and a closely related product's page. Consumers would click on an RP on the FP's page because of its easy navigation and relevance. Thus, it is difficult to tease out the pure navigation effect from the relevance effect of recommendations. One way to uncover the navigation effect is to compare the effects of recommendations on the view and purchase probability of RPs with high- (strongly related) versus low-affinity scores (weakly related) with the FP. If pure navigation effect is the primary mechanism that drives the benefits of recommendations to consumers, we should expect a similar view and purchase probability between weakly and strongly related FP-RP pairs. Accordingly, we estimate the following specifications at the FP-RP pair level:

$$Y_{frst} = \alpha_f + \alpha_t + \beta_1 \times Rec_{frst} + \varepsilon_{frst} \; ; \tag{8.1}$$

$$Y_{frst} = \alpha_f + \alpha_t + \beta_1 \times Rec_{frst} + \beta_2 \times High_{frt} + \beta_3 \times High_{frt} \times Rec_{frst} + \varepsilon_{frst} \; ; \tag{8.2}$$

where *f* denotes FP, *r* denotes RP, *s* denotes sessions, and *t* denotes days. α_f and α_t, respectively, denote the FP- and day-fixed effects. The dependent variable Y_{frst} could be $RP\ View_{frst}$ or $RP\ Purchase_{frst}$, denoting whether the RP *r* is viewed and purchased, respectively, after FP *f*'s page view in session *s* in day *t*. Rec_{frst} is an indicator for the treated sessions. We estimate the specification (8.1) on two distinct datasets: (i) weakly related FP-RP pairs with affinity scores in the bottom 25 percentile of the affinity score distribution, and (ii) strongly related FP-RP pairs with affinity scores in the top 25 percentile of the affinity score distribution. $High_{frt}$ is an indicator variable for the FP-RP pair for strongly related RPs.

Table 11 reports the results. We find a significantly higher positive effect of recommendations for strongly related RPs than weakly related RPs (Columns (2) vs. (1), and Columns (5) vs. (4)). We find similar results in the pooled regression on the data on weakly and strongly related FP-RP pairs in Columns

(3) and (6). While the navigation effect on RP view/purchase probabilities should be same for weakly and strongly related RP-FP pairs, significantly higher view/purchase probabilities for strongly related RPs indicate the pure navigation effect is less likely to be a primary underlying mechanism that drives the benefits of recommendations to visitors.

Table 11. Effects of Recommendations for Weakly vs. Strongly Related RPs

	DV: Whether RP viewed			DV: Whether RP purchased		
	(1) Weakly Related	(2) Strongly Related	(3) Pooled	(4) Weakly Related	(5) Strongly Related	(6) Pooled
Rec. Indicator	0.013*** (0.000)	0.041*** (0.001)	0.013*** (0.000)	0.0004*** (0.0001)	0.001*** (0.000)	0.0004* (0.000)
High			0.032*** (0.001)			0.004*** (0.000)
Rec. Indicator ×High			0.027*** (0.001)			0.001*** (0.000)
R^2	0.0384	0.0317	0.0427	0.0239	0.0102	0.0130
Obs.	993,922	992,450	1,986,372	993,922	992,450	1,986,372

Notes: + p < 0.10, * p < 0.05, ** p < 0.01, *** p < 0.001

6.2.5.2 Information Effect

The benefits of recommendations could be due to showing more products (not necessarily algorithmically identified related products) on the FPs' pages. Thus, the effect of recommendations may be due to the availability of more product information under recommendations.

First, we note that both product category pages and search result pages contain several thumbnail-sized product pictures with product information such as product prices. Usually, these pages display more products than the four RPs shown under recommendations on the FP's page. Appendix B shows an example of the women's dress category main page and the search result page for the keyword search of "Floral women's top" on the retailer's website. These pages contain information on over 12 products compared to similar information on four RPs on the FP's page under recommendations. This fact suggests that control visitors with more product category pages and search result page views could see information on more products on these two types of pages. Thus, more product information on the FP pages alone is not likely to explain why treated visitors have a higher purchase probability than control visitors.

Second, Section 5.1 shows that visitors are more likely to view and purchase RPs with high-affinity scores with the FP. We further show that visitors are more likely to buy those viewed products in the purchase funnels with higher affinity scores with the first viewed product. These empirical results show that visitors are more likely to react to relevant product information, but not all product information. In Section 6.1, we find that displaying recommendations in the treated purchase funnels increases the average affinity score of the products browsed in the purchase funnels. After that, we use the instrumental variable (IV) approach to estimate the effect of the average affinity score of the browsed products on purchase probability due to the exogenous availability of recommendations. As reported in column (2) of Table 7, we find that a higher average affinity score of the browsed products could lead to a higher purchase probability. Overall, we find that it is the relevance of RPs with the FP (captured by the affinity scores) rather than the amount of (content summarization) information about the RPs on the FP's page that drives our results.

Lastly, results on weakly versus strongly related RPs in Section 6.2.5.1 suggest that showing relevant (strongly related) RPs on the FP pages benefits RP sales but the effect for weakly related products (still more information) is much smaller. Overall, more information about RPs on the FP pages is unlikely to be the primary underlying mechanism that drives the benefits of recommendations to consumers.

6.3. Substitution of Existing Search Tools with Recommendations

The retailer's website offers visitors two additional search tools besides product recommendations: (i) product category-based search through the hierarchical organization of products on the website and (ii) keyword-based search. We measure visitors' product category-based and keyword-based searches with the number of product category/subcategory page views and search result page views in a purchase funnel. Visitors endogenously decide which and how much of a search tool to use. In our experiment, we additionally provide the recommendation tool to some randomly selected visitors. Thus, our experimental setup allows us to examine the effect of the exogenous availability of recommendations on visitors' endogenous choice of search tools on the retailer's website.

We estimate the effect of the recommendation availability and usage (measured by the number of

RP page views) on the usage of existing search tools on the website with the following specifications:

$$Y_i = \beta_c + \beta_1 \times Rec_i + \varepsilon_i; \tag{9.1}$$

$$Y_i = \beta_c + \beta_1 \times NRPView_i + \varepsilon_i; \tag{9.2}$$

where the dependent variable Y_i denotes the usage of two existing search tools on the website, i.e., the number of search result page views or product category/subcategory page views in the purchase funnel i. All other variables have the same meanings as in the previous specifications.

Table 12 reports the estimated coefficients from specifications (9.1) and (9.2). In columns (1) and (4), we find negative and significant coefficients for the recommendation indicator variable in the specification (9.1). These estimates indicate that visitors reduce their keyword-based and product-category-based searches with the availability of recommendations. Specifically, visitors with recommendations view 0.02 fewer search result pages and 0.10 fewer product subcategory pages in their purchase funnels.

Table 12. Effects of Recommendations on Usages of Other Search Tools

DV:	*No. of Search Result Page Views*			*No. of Product Category Page Views*		
	(1)	(2)	(3)	(4)	(5)	(6)
Rec. Indicator	-0.02^{**} (0.01)			-0.10^{***} (0.02)		
No. of RP Page Views		0.57^{***} (0.05)	-0.12^{***} (0.03)		1.85^{***} (0.29)	-0.65^{***} (0.10)
Instrumental Variable		No	Rec. Indicator		No	Rec. Indicator
Obs.	573,665	573,665	573,665	573,665	573,665	573,665

Notes: + p < .10, * p < .05, ** p < .01, *** p < .001. Standard errors cluster corrected at product subcategory level in parentheses.

We suspect that the number of RP page views ($NRPView_i$) in the specification (9.2) would be endogenous, as visitor-level unobserved characteristics may determine both their usage of recommendations and other search tools on the retailer's website. Columns (2) and (5) of Table 12 show a positive and significant correlation of the number of RP page views with the number of the search result page and product category page views, respectively, from a fixed effect OLS estimation of the specification (9.2). However, the estimate of a positive correlation is biased due to the omission of visitor-level factors.

To address the endogeneity of $NRPView_i$ in the specification (9.2), we instrument it with the indicator variable of recommendations. In the previous sections, we have already shown that the recommendation indicator satisfies the conditions of exogeneity and relevance for instrumental variables. Thus, we find the effect of the number of RP page views attributed to the availability of recommendations on the number of search result page views. Columns (3) and (6) in Table 12 report the estimates from the 2SLS regressions. The results show negative and significant coefficients for the number of RP page views from the specification (9.2), suggesting the substitution between the RP page views and the number of the search result page (and product subcategory page views).

7. Conclusion

We estimate the benefits of product recommendations to visitors from a randomized field experiment on the website of a US apparel and home goods retailer. We impute a product's net value to consumers and fit with their tastes based on unique data on the affinity scores computed by the recommendation algorithm. We find that product recommendations help consumers search for higher-value products that are lower-priced, fit their tastes better, or both. The results also show that the discovery of higher-value products results in a higher purchase probability under recommendations. Finally, when made available, consumers substitute their usage of other search tools on the website with product recommendations.

7.1. Managerial Implications

Our findings have important practical implications for online retailers, policymakers, regulators, and item-based CF recommender systems design.

Implications for online retailers. Our findings show that item-based CF recommender systems help consumers find products of higher net value and better horizontal fit. Given these benefits of recommendations, online retailers offering product recommendations could expect a higher level of consumer satisfaction and retention than those that do not. Our findings that item-based CF recommender systems substitute other search tools on the website inform online retailers to allocate more resources to develop product recommendations. Finally, our moderation analysis guides online retailers under what conditions product recommendations could bring higher benefits to consumers. Specifically, consumers may enjoy higher price savings in product categories with higher price dispersions and higher horizontal fit (such as women's products). Online retailers selling such products could garner higher customer satisfaction and retention by providing recommendations.

Implications for policymakers and regulators. Policymakers are concerned that online retailers use recommender systems to generate higher sales at the expense of consumer welfare. Such concerns have given rise to significant regulations constraining online retailers from using recommender systems. For example, the recent EU's Digital Services Act (ARTICLE 19) mandates that online retailers ensure transparency and diversity of exposure.[13] Our item-based CF recommender system is purely based on consumers' product co-views and co-purchases. It does not include the retailer's strategic interests, such as profit maximization or promoting specific products (e.g., high-margin or lower-selling products). We show that compared to consumers shopping without recommendations, pure item-based CF recommender systems could help consumers find and buy products of higher value and better horizontal taste fit. Thus, our findings alleviate the concern of regulators and policymakers on transparency and diversity of exposure with item-based CF recommender systems.

Consumer privacy concerns have also led to regulations on the recommendation systems used by online retailers. For example, the EU's General Data Protection Regulation (GDPR) prevents online retailers from using cookies without consumers' consent, potentially limiting the ability of recommender systems.[14] We note that item-based CF recommender systems, like the one used in our field experiment, utilize only consumer browsing and purchasing data without any identity information. Our paper shows that only using aggregate consumer browsing and buying data by recommender systems could bring substantial consumer benefits. Regulators and policymakers may pay more attention to such privacy-protected recommender systems.

Implications for recommendation design. Finally, our paper introduces a novel method of

[13] Source: https://www.article19.org/resources/eu-regulation-of-recommender-systems-in-the-digital-services-act/#:~:text=Recommender%20systems%20select%20content%20based,that%20engages%20users%20the%20most

[14] Source: https://www.forbes.com/sites/andrewarnold/2018/05/07/how-gdpr-and-changing-legislation-will-impact-digital-advertising/?sh=1b6c68892450

inferring a product's net value and horizontal fit to consumers from the computed scores of item-based CF algorithms. Therefore, our approach can reveal the total value of different RPs on an FP's page in terms of horizontal fit and price. More nuanced item-based CF recommendation algorithms could be designed based on this new information. For example, future item-based CF algorithms may explore attaching a higher weight to a product's horizontal fit while recommending products in product categories (or consumer populations) with high taste heterogeneity. We hope our method will encourage future research to develop more nuanced recommendation algorithms.

7.2. Generalizability of Findings

Our findings are for the item-based CF algorithm, which is among the most widely used recommendation algorithm in practice.[15] Our results would generalize to the broad category of item-based CF algorithms computing similarity between products based on various criteria such as their co-views and co-purchases or consumer ratings (Sarwar et al. 2001, Linden et al. 2003, Kumar and Hosanagar 2019, Lee and Hosanagar 2019, 2021, Li et al. 2022). Second, our findings are for online apparel and accessories, the second highest selling product category in E-commerce. Thus, our findings, applicable to over US $185 billion in E-commerce apparel sales worldwide, are academically and managerially significant.[16] Third, our findings are for substitute product recommendations. We conducted an extensive survey and found that most E-commerce apparel retailers recommend substitute products. Moreover, our approach to computing a product's net value and horizontal fit would also apply to complementary product recommendations. Our approach should motivate future research to examine the underlying reasons for the benefit of complementary recommendations to consumers.

7.3. Limitations and Future Research Directions

Our research has several limitations that provide opportunities for future research. Our estimates are for a widely used item-based CF recommendation system in the context of apparel and accessories. Our findings

[15] Source: https://www.grandviewresearch.com/industry-analysis/recommendation-engine-market-report

[16] Source: https://www.emarketer.com/content/us-ecommerce-by-category-2021

may not be generalizable to other categories of recommender systems, such as those studied by Jiang et al. (2015) and Ghose et al. (2012). Examining the benefits of other recommendation systems to consumers could be an interesting avenue for future research. Another limitation of our experiment is the possible imprecise identification of visitors. A treated visitor may be misclassified as a control visitor when she uses a different device. However, this is a common limitation of most online experiments. Such an overlap of visitors to different treatment conditions only makes our results more conservative. Future research should devise better methods of visitor identification. Finally, our approach to computing a product's net value and horizontal fit for item-based CF substitute product recommendations would also apply to complementary and supplementary product recommendations. We hope future research will utilize our approach to estimate the benefit of supplementary/complementary recommendations to consumers.

Acknowledgments

The authors gratefully acknowledge generous research support from Hi! PARIS Fellowship, HEC Foundation, and the Public Utility Research Center (PURC) of the University of Florida. The authors thank the department editor Professor D.J. Wu, the associate editor, and the three anonymous reviewers for their constructive and insightful suggestions. The autors greatly appreciate the valuable comments obtained while presenting the research at CIST 2021, ZEW Conference on the Economics of Information and Communication Technologies 2021, CODE@MIT 2020, WISE 2020, Temple University, George Mason University, University of Texas at Dallas, and Tel Aviv University. The authors also thank Professor Sumit Sarkar (University of Texas at Dallas) and Professor Tianxin Zou (University of Florida) for their helpful comments.

References

Abdollahpouri H, Adomavicius G, Burke R, Guy I, Jannach D, Kamishima T, Krasnodebski J, Pizzato L (2020) Multistakeholder recommendation: Survey and research directions. *User Modeling and User-Adapted Interaction.* 30(1) 127-158.

Adomavicius G, Bockstedt J, Curley S, Zhang J (2019) Reducing recommender systems biases: An investigation of rating display designs. *MIS Quarterly*. 43(4): 1321-1341.

Ali A, Zhang W (2015) CEO tenure and earnings management. *Journal of Accounting and Economics.* 59(1) 60-79.

Basu S, Fernald JG, Kimball MS (2006) Are technology improvements contractionary? *American Economic Review.* 96(5) 1418-1448.

Bleier A, Eisenbeiss M (2015) Personalized online advertising effectiveness: The interplay of what, when, and where. *Marketing Science.* 34(5) 669-688.

Bound J, Jaeger DA, Baker RM (1995) Problems with instrumental variables estimation when the correlation between the instruments and the endogenous explanatory variable is weak. *Journal of the American Statistical Association.* 90(430) 443-450.

Brynjolfsson E, Rock D, Syverson C (2021) The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics.* 13(1) 333-72.

Burtch G, Ghose A, Wattal S (2013) An empirical examination of the antecedents and consequences of contribution patterns in crowd-funded markets. *Information Systems Research.* 24(3) 499-519.

Chapelle O, Zhang Y (2009) A dynamic bayesian network click model for web search ranking. *Proceedings of the 18th International Conference on World Wide Web.* 1-10.

Chen Y, Yao S (2017) Sequential search with refinement: Model and application with click-stream data. *Management Science.* 63(12) 4345-4365.

De P, Hu Y, Rahman MS (2010) Technology usage and online sales: An empirical study. *Management Science.* 56(11) 1930-1945.

Dellaert BG, Häubl G (2012) Searching in choice mode: consumer decision processes in product search with recommendations. *Journal of Marketing Research.* 49(2) 277-288.

Dinkelman T (2011) The effects of rural electrification on employment: New evidence from South Africa. *American Economic Review.* 101(7) 3078-3108.

Dou Y, Hope O-K, Thomas WB (2013) Relationship-specificity, contract enforceability, and income smoothing. *The Accounting Review.* 88(5) 1629-1656.

Ghose A, Ipeirotis PG, Li B (2012) Designing ranking systems for hotels on travel search engines by mining user-generated and crowdsourced content. *Marketing Science.* 31(3) 493-520.

Ghose A, Ipeirotis PG, Li B (2019) Modeling consumer footprints on search engines: An interplay with social media. *Management Science.* 65(3) 1363-1385.

Goldenberg J, Oestreicher-Singer G, Reichman S (2012) The quest for content: How user-generated links can facilitate online exploration. *Journal of Marketing Research.* 49(4) 452-468.

Hadi AS, Chatterjee S (2015) *Regression analysis by example* (John Wiley & Sons). 250.

Honka E (2014) Quantifying search and switching costs in the US auto insurance industry. *The RAND Journal of Economics.* 45(4) 847-884.

Jiang C, Duan R, Jain HK, Liu S, Liang K (2015) Hybrid collaborative filtering for high-involvement products: A solution to opinion sparsity and dynamics. *Decision Support Systems.* 79 195-208.

Jones JJ (1991) Earnings management during import relief investigations. *Journal of Accounting Research.* 29(2) 193-228.

Jorgenson DW, Stiroh KJ (1999) Information technology and growth. *American Economic Review.* 89(2)

109-115.

Kim K, Mauldin E, Patro S (2014) Outside directors and board advising and monitoring performance. *Journal of Accounting and Economics.* 57(2-3) 110-131.

Kumar A, Hosanagar K (2019) Measuring the value of recommendation links on product demand. *Information Systems Research.* 30(3) 819-838.

Kumar A, Tan Y (2015) The demand effects of joint product advertising in online videos. *Management Science.* 61(8) 1921-1937.

Kwark Y, Chen J, Raghunathan S (2017) Platform or wholesale? A strategic tool for online retailers to benefit from third-party information. *MIS Quarterly.* 41(3) 763-785.

Lambrecht A, Tucker C (2013) When does retargeting work? Information specificity in online advertising. *Journal of Marketing Research.* 50(5) 561-576.

Lee D, Hosanagar K (2019) How do recommender systems affect sales diversity? A cross-category investigation via randomized field experiment. *Information Systems Research.* 30(1) 239-259.

Lee D, Hosanagar K (2021) How do product attributes moderate the impact of recommender systems? *Management Science.* 67(1) 524-546.

Li X, Grahl J, Hinz O (2022) How do recommender systems lead to consumer purchases? A causal mediation analysis of a field experiment. *Information Systems Research.* 33(2) 620-637.

Lin Z, Goh KY, Heng CS (2017) The demand effects of product recommendation networks: An empirical analysis of network diversity and stability. *MIS Quarterly.* 41(2) 397-426.

Linden G, Smith B, York J (2003) Amazon. com recommendations: Item-to-item collaborative filtering. *IEEE Internet Computing.* 7(1) 76-80.

Ngwe D, Ferreira KJ, Teixeira T (2019) The impact of increasing search frictions on online shopping behavior: Evidence from a field experiment. *Journal of Marketing Research.* 56(6) 944-959.

Oestreicher-Singer G, Sundararajan A (2012) The visible hand? Demand effects of recommendation networks in electronic markets. *Management Science.* 58(11) 1963-1981.

Ratchford BT, Srinivasan N (1993) An empirical investigation of returns to search. *Marketing Science.* 12(1) 73-87.

Reinganum JF (1982) Strategic search theory. *International Economic Review.* 23(1) 1-17.

Sarwar BM, Karypis G, Konstan JA, Riedl J (2001) Item-based collaborative filtering recommendation algorithms. *In Proceedings of the 10th international conference on World Wide Web.* 285-295.

Seiler S, Pinna F (2017) Estimating search benefits from path-tracking data: Measurement and determinants. *Marketing Science.* 36(4) 565-589.

Singh PV, Sahoo N, Mukhopadhyay T (2014) How to attract and retain readers in enterprise blogging? *Information Systems Research.* 25(1) 35-52.

Solow RM (1956) A contribution to the theory of economic growth. *The Quarterly Journal of Economics.* 70(1) 65-94.

Solow RM (1957) Technical change and the aggregate production function. *The Review of Economics and Statistics.* 39(3) 312-320.

Sutton J (1986) Vertical product differentiation: some basic themes. *The American Economic Review.* 76(2) 393-398.

Thorat PB, Goudar R, Barve S (2015) Survey on collaborative filtering, content-based filtering and hybrid recommendation system. *International Journal of Computer Applications.* 110(4) 31-36.

Ursu RM, Wang Q, Chintagunta PK (2020) Search duration. *Marketing Science.* 39(5) 849-871.

Weitzman ML (1979) Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*. 47(3) 641-654.

Zhang X, Ferreira P, Godinho de Matos M, Belo R (2021) Welfare properties of profit maximizing recommender systems: Theory and results from a randomized experiment. *MIS Quarterly.* 45(1) 1-43.