Measuring the Value of Recommendation Links on Product Demand

Anuj Kumar
Warrington College of Business
University of Florida
(akumar1@ufl.edu)

Kartik Hosanagar
The Wharton School
University of Pennsylvania
(kartikh@wharton.upenn.edu)

This Version
(June 2018)

Abstract

Recommending substitute products on focal products' pages on an e-commerce website can impact product sales in two ways. First, the visibility of a product as a recommendation on other products' pages may increase its exposure and result in a greater number of its page views. Second, visibility of substitute products on the product's page may cannibalize its own sales while resulting in greater exposure for the substitute products. The net impact of these opposing effects is unclear. We conduct a randomized experiment on a fashion apparel retailer's website to answer the following questions: (1) What is the causal value of recommendation links from a product to its recommended products in terms of the additional sales for both the product and its recommended products, and (2) how does the value of a product's recommendation links vary based on its network characteristics, such as its PageRank and the strength of its relationship with neighboring products. We find that due to recommendation, on average: (1) daily number of product page views increased by 7.5 percent and (2) conditional on a product's page view, its sales decreased by 1.9 percent and the sales of its recommended substitutes increased by nine percent. On average, recommendation links of a product result in an 11 percent gain in total sales of the product and its recommended substitutes. However, these gains are not evenly distributed among all products. We find that while number of page views for a product are positively affected by the number and strength of its incoming links, its sales (its recommended products' sales) conditional on its page view are negatively (positively) affected by the strength of its outgoing links. We conduct policy simulations to highlight how retailers and producers can apply this knowledge by engineering the recommendation network through sponsored links.

Keywords: *Product recommendation network; electronic commerce; randomized field experiment; sponsored product advertising; average treatment effect.*

1.0 Introduction

Recommendation systems are widely used by major online retailers including Amazon, Walmart and iTunes. Once a consumer reveals her preference for a product category/type by visiting a product's page, recommending closely related products on the product's page can reduce customers' search costs, especially when consumers have to evaluate several thousand products on a retailer's website. Recent research has established that these recommendations influence consumer choice (Senecal and Nantel 2004) and generate a lift in sales (De et al. 2010).

Algorithmic recommendations on product pages such as "people who viewed this product also viewed these other products" create links between products resulting in a network of interconnected products on a retailer's website, often referred to as the product recommendation network. In this network, nodes represent the products on a retailer's website and directed edges represent recommendation links from one product's page to another product's page. The recommendation network creates alternative browsing paths for consumers and affects their browsing and purchase behaviors. From a retailer's standpoint, there is emerging interest in understanding at a granular level the impact of these links on product sales (Oestricher-Singer and Sundararajan 2012; Lin et al. 2017). At a product level, recommendation links can impact a product's sales in two ways. First, the visibility of a product as a recommendation on other products' pages may increase its exposure and result in a greater number of its page views. Second, visibility of close substitutes (complements) on a product's page may cannibalize (complement) its own sales while resulting in greater exposure for the substitutes (complements). The net impact of these effects is unclear especially in the context of recommendation engines that recommend substitute products as opposed to complements.

Recent studies have investigated demand spillovers between products due to recommendation links (Oestricher-Singer and Sundararajan 2012), the "network value" of a product in terms of the sales it generates for other products through outgoing recommendation links

(Oestricher-Singer et al. 2013), and the impact of product category diversity of incoming recommendation links on the demand of a focal product (Lin et al 2017). These studies have demonstrated the economic significance of links embedded within a recommended network. However, existing studies have not focused on how the economic value of a recommendation link varies based on the source and strength of the links. In terms of link source, different products are located in different parts of the recommendation network which, in turn, might impact the value of a link from their product pages. In terms of link strength, recommended products may sometimes be closely related or weakly related to a focal product. We study the impact of both of these factors on the value of a recommendation link between products.

Our work also extends prior work by decomposing the impact of recommendation links in terms of number of views a product receives versus its sales conditional on being viewed. This decomposition is particularly important when substitute products are recommended. While incoming recommendation links help increase a product's views, outgoing links to substitute products may exert a negative impact on purchase probability. These managerially significant interactions are masked in an aggregate analysis of sales in the previous studies (Oestricher-Singer and Sundararajan 2012, Lin et al. 2017). Decomposing these two opposing interactions in the context of recommendation of substitute products allows us to better unlock the mechanism through which recommendation networks impact product success.

Estimation of the economic value of a recommendation link and how it varies based on the above factors is complicated by endogeneity concerns. The most important challenge is simultaneity – because popular recommendation engines are based on past views and purchases, the recommendation links are driven by product popularity just as sales may be driven by the links. Further, any effort to associate variation in links between products to variation in sales has to clearly account for the fact that purchases for substitute products are likely correlated in the first place.

To address these challenges, we partnered with a large online retailer to implement a randomized field experiment in which a random sample of visitors on the retailer's website in a treatment group observe product recommendations (i.e., see substitute products on a product's page) while visitors in a control group do not see any recommendations. We collect visitor's session-level data to compute the page views and purchase probabilities for each product in a world with recommendations versus one without recommendations and attribute the difference to the impact of recommendation links. This approach allows us to answer in a more granular manner the following questions: (1) What is the causal value of recommendation links from a product to its recommended products on a retailer's website in terms of the additional sales for both the product and its recommended products, and (2) how does the value of a product's recommendation links vary based on its network characteristics, such as its PageRank (a measure of connectedness of the product in the recommendation network) and the strength of its relationship with neighboring products (computed based on their coviews and copurchases).

We find that the recommendation links of a product, on average, result in an 11 percent lift in total sales of the product and its recommended substitutes. We further decompose these aggregate estimates into component estimates, namely the impact on product page views and product sales conditional on the product page view. We find that, on average, the daily number of product page views ("exposure") increased by 7.5 percent due to recommendations. Further, conditional on a product's page view, its sales decreased by 1.9 percent (due to the "distraction" of outgoing recommendation links to its substitutes) whereas the sales of its recommended substitutes increased by nine percent due to recommendations.

We further find that the value of a product's recommendation link varies with the network characteristics of the product. While the number of product page views are positively affected by the number and strength of its incoming links, the products' sales (its recommended products' sales) conditional on its page view are negatively (positively) affected with the strength of its outgoing links. The value of a product's recommendation links could come from merely exposing other products (exposure effect) and/or from showing related products (recommendation effect) on a product's page. We find that the value of a product's recommendation links in our field setup came from the recommendation effect but not from the exposure effect. Our analysis, therefore, identifies a hitherto underexplored, yet highly significant, driver of product performance in online retail, namely its network characteristics in the product recommendation network.

Our results inform managerial practice for retailers. While recommending substitute products can cannibalize sales, we demonstrate that the net impact of recommending substitute products is positive. This is because lower search costs appear to drive a significant increase in purchase incidences. Our results also demonstrate why recommendations need to be viewed as a strategic tool to drive product sales and market share. Thus, just like firms expend effort on search engine optimization (SEO) – for example by trying to obtain hyperlinks from other popular websites – to increase website traffic, we advocate that they need to analyze recommendation networks and invest in driving co-views and co-purchases with relevant popular products. This will help enhance product discovery. Similarly, sellers can benefit by advertising their products on other product pages via ad auctions run by retailers such as Amazon. Finally, we conduct policy simulations to investigate the potential sales gains to producers from sponsored advertising of their products with related products on the retailer's website. In turn, this informs their decision on how much to bid for a sponsored listing.

The rest of the paper is organized as follows. In the next section, we discuss the related literature. In Section 3, we describe our research setting and experimental design. We present our data in Section 4, empirical models, results, and robustness checks in Section 5, and policy simulations in section 6. In Section 7, we conclude with managerial implications of our research and outline future research directions.

2.0 Literature Review

There is a vast body of work on the design of recommendation engines (see Adomavicious et al. 2008 for an overview). In terms of their impact on users and firms, many recent studies have shown their positive impact on individual consumer choice (Senecal and Nantel 2004) as well as overall sales generated by retailers (De et al. 2010). Recent research has also investigated their impact on sales diversity (Fleder and Hosanagar 2009, Lee and Hosanagar 2018).

An emerging stream of work has gone beyond the aggregate impact of recommendations to highlight the importance of product networks created by recommendation engines. Oestricher-Singer and Sundararajan (2012) show that visibility of co-purchase links between books significantly increases the correlations in their sales. Lin et al. (2017) further analyze the effect of the diversity and stability of links between focal and recommended products on their demand. Oestricher-Singer et al. (2013) suggest that products generate value to other products through outgoing recommendation links and derive value from incoming links. Using the product network on Amazon, they decompose a product's sales into its "intrinsic" sales and the sales attributable to incoming recommendation links from other product pages. They also compute a product's network value based on the sales it generates for other products through outgoing recommendation links. One of the challenges with these studies that are based on observational data is that they lack the contrast needed to separate the effect of recommendations on views/sales from the effect of views/sales on recommendations. As Lin et al. (2017) comment on the limitations of their methods as well as those of other studies based on observational data: "these approaches may not have fully controlled for all potential sources of endogeneity bias. We thus do not make absolute causality claims." To address these challenges, Sharma et al. (2015) exploit demand shocks on certain products to estimate the effect of recommendation systems on the page views of recommended products. Similarly, Carmi et al. (2017) show that a demand shock to a book due to its appearance on the Oprah Winfrey show affects the sales of the book's immediate neighbor in the recommendation network as well as that of its second and third degree neighbors.

The above stream of work has clearly shown that dyadic links between related products affect their sales. While it is now clear that recommendation links between products carry economic value, much work is needed for us to understand what factors affect the value of these links and in which direction. Poor understanding of the drivers of value in recommendation networks has also meant that practitioners are often uncertain about how to exploit recommendation networks to drive business performance.

One framework that can help us understand the factors that impact the value from dyadic links between products in a network comes from the study of social networks. In social networks, research has shown that the value of a tie between two entities depends on its location in the network and the strength of the tie. Social network theories posit that central location in a network accord higher influence and access to information to the entities (Podolny 2001). Further, the strength of ties – strong versus weak – affect the value of the ties (Krackhardt et al. 2003, Granovetter 1973). Granovetter (1973) emphasized the importance of weak ties as bridges to other parts of the network. As an example, he showed that workers often found job opportunities from their weak ties which offered new information. In contrast, Burt (1992) argued that bridging ties could be more beneficial if they are strong. On similar lines, other studies have shown that stronger, more intensive, and long-term ties between entities in social networks offer the greatest value (Dore 1983, Uzzi 1996).

Motivated by this literature, our study extends the emerging literature on product networks by asking three important questions:

1. How does the *location of products* in a product network and *strength of ties* between them affect their demand? Just like in social networks, the location of a product in the network and the strength of its ties with its neighbors should determine its visibility and hence its demand. While previous studies on product network have considered the effect of location of products in their network (such as its indegree or PageRank), none of them have examined how the strength of their ties with neighboring products affects their demands.

Understanding the impact of link strength helps unmask how much of the impact of a recommendation stems from the identification of a *relevant* product versus the mere exposure of another product from a retailer's catalog or the recommendation of a popular product that most consumers like. This, in turn, helps clarify how much the quality of the recommendation algorithm matters relative to a simple exposure effect wherein consumers are exposed to more products.

2. How does recommendation of substitute products affect the number of views and purchase probability of products? Prior studies often look at the impact of recommendations on book pages wherein recommendation links may direct the consumer to either a complementary product or a substitute product. For example, a recommendation on the page for a book on wine may be for a book on cheese (complementary product) or a different book on wine (substitute product). Thus, these studies do not distinguish between substitute and complementary ties between focal and recommended products. The positive effect of recommending strictly complementary products in online videos has been demonstrated by Kumar and Tan (2015), but to our knowledge, no prior study has examined the demand effects of recommending strictly substitute products. Thus, although the practice of recommending substitute products is common (Macy's.com, JCPenny.com, and other e-tailers), extant research offers almost no guidance on how substitutive ties between neighboring products affect their demand. The case of substitute products is particularly interesting because incoming links help increase a product's views but outgoing recommendations links to substitute products may reduce its purchase probability conditional on view. Therefore, while it is not surprising that recommendations to complementary products can help increase sales, the net impact of recommendations on substitute products is not readily obvious.

3. How can retailers and advertisers think strategically about recommendation links and extract value from recommendations on a retailer's website? Recognizing that

recommendation links carry economic value, many prominent e-retailers, including Amazon.com and Alibaba.com, have started to show sponsored product recommendations on product pages on their website. We analyze how a website owner (retailer) can create a market for producers/advertisers to bid for sponsored display of their products on other product pages and how advertisers, in turn, can compute the additional sales garnered by a product from its recommendation on other products' pages. While many papers have shown that recommendation links carry economic value, no paper we have seen connects that value to managerial action for a producer.

We examine these questions based on a randomized field experiment with direct observations of the recommendation network as well as data on product views and sales, thereby enabling clean identification of both the direction and size of the effect of a product's recommendation link. To the best of our knowledge, this is the first study to estimate the causal economic value of a product's recommendation links in online retail. We focus exclusively on the effect of recommending substitute product which, as pointed out above, are both commonly used on retail websites and have economically unclear impact on sales. Unlike past studies, we decompose the impact on a product's sales into the impact on its views and purchase probability (conditional on views). Finally, we provide managerial guidance on how retailers and advertisers can utilize our findings to compute the additional sales generated from an additional link on a product's page. These calculations allow a producer to think strategically about sponsored recommendation links on a retailer's website much like they might analyze and value sponsored placement in search results.

3.0 Field Setup

We conducted a field study on the website of a mid-sized fashion retailer in the US.[1] The retailer has annual revenues of over US $400 million and sells its products via its physical stores and website. We examine the retailer's online sales in this paper, which accounts for over 10 percent of its total sales.

3.1 Website Organization

The retailer offers over 35,000 products for sales on its website. These products are organized under different product categories such as women, men, kids, juniors, and home goods on the retailer's website. The main page of a product category can be reached by clicking on its hyperlinks on the home page of the website. The main pages of product categories provide hyperlinks for the product subcategories under them. For example, the main page for women's products provides hyperlinks to subcategories such as tops, shorts, skirts, dresses, and pants. Clicking the hyperlink for a product subcategory takes visitors to the product subcategory's main page, where 24 to 30 thumbnail-sized pictures of products within that subcategory are displayed. Additional products within the subcategory can be viewed by navigating to the next page. For example, women's tops are displayed on over 20 pages on the retailer's website. Clicking the thumbnail picture of a product takes visitors to the product description page (or simply, a product page). An example of a product page for a women's top is shown in Figure 1.

A product page displays an enlarged picture of the product along with the product description, retail price, offered discount, average customer review ratings, and available colors and sizes. The product on its product page is referred to as the focal product. Thus, "Cathy Daniels Womens Embellished Palm Tree Top" is the focal product in Figure 1. The product page also displays pictures of four related products, referred to as recommended products under the heading "MORE OPTIONS". For example, pictures of four women tops, shown with the focal top, are the

[1] The retailer's identity is not divulged due to the Non-Disclosure Agreement

recommended products in Figure 1. The retail prices, offered discounts, and average customer review ratings for these recommended products are also displayed on the focal product's page. Clicking on a recommended product's picture takes the visitor to its product page where it is now the focal product for the user.

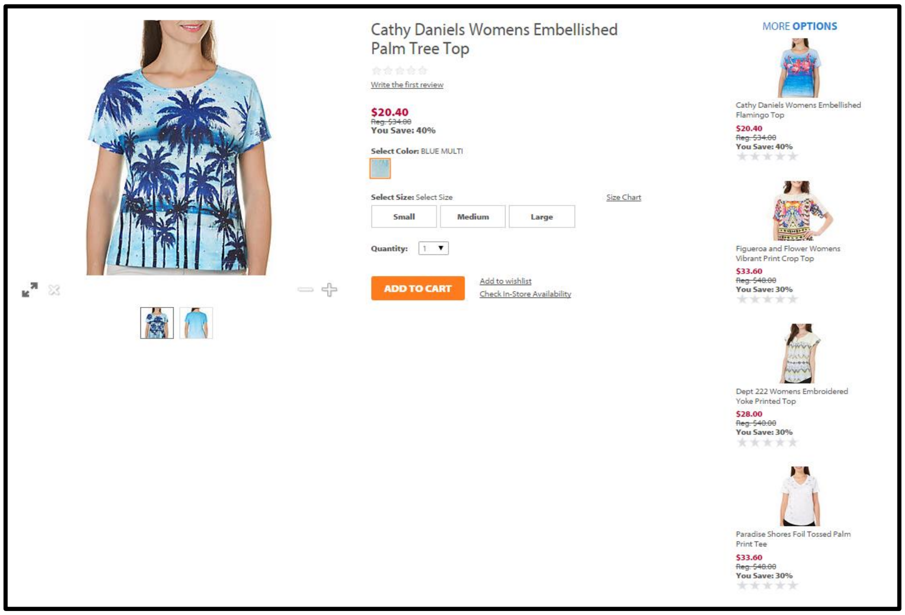

Figure 1: Product page

The recommended products to be displayed on a focal product's page are chosen on the basis of their affinity scores with the focal product, which is computed by the IBM Coremetrics recommendation engine. For a focal-recommended product pair, the affinity score is based on the following four values computed over the last 30 days: (1) number of sessions in which the focal and recommended products are viewed together in the same session (called view-to-view score); (2) number of sessions in which the focal product is viewed but the recommended product is bought in the same session (called view-to-buy score); (3) number of times the focal and recommended products are bought by the same person not necessarily in the same session (called buy-to-buy score); and (4) number of sessions in which the focal product is abandoned and recommended product is bought in the same session (called abandon-to-buy score). The affinity score for a focal-recommended product pair is the weighted sum of these four scores. A common practice is to weigh the view data heavily when the retailer would like to show substitute products and the purchase data heavily when the retailer seeks to show complementary products. The formula used by our retail partner, which is in fact the default setting on Coremetrics for recommending substitutes, is: affinity score = 70* view-to-view score + 20 * view-to-buy score + 5 * buy-to-buy score + 5 * abandon-to-buy score, wherein the highest weight is applied to the view data in an effort to show substitute products. To further ensure that only substitute products are recommended on a focal product's page, the product recommendation engine additionally applies the business rule that only products in the same product subcategory as the focal product can be recommended. Accordingly, the affinity scores of each product with all other products on the retailer's website are computed daily at 4 AM and data on the top 15 related products based on their affinity scores with each product is stored.[2] The top four related products in the same product subcategory and having the highest affinity scores with a product are recommended on its product page.

3.2 Search and Navigation Options on the Website

As indicated in the previous section, the retailer organizes over 35,000 products on its website into different product categories and into different product subcategories within each product category. A visitor can navigate back and forth to search for products under different categories by going over to the home page, category main page, subcategory main page, and the product pages on the retailer's website. For example, a visitor's search path could be home page → Women's category main page → Women's Tops subcategory main page → A specific Women's Top's page. Besides navigating to a desired product's page from its product category and subcategory main pages, a visitor can directly search for the desired product or category of products by typing in the text search bar available on each page of the retailer's website. The search bar allows visitors to do search by keywords, item number and page code. For example, typing "women's summer dress"

[2] The affinity scores for most products become zero before their 15th ranked related products. So, not all products have 15 related products with non-zero affinity scores. For products that have non-zero affinity scores up to their 15th ranked related product, the affinity scores beyond the 15th ranked related product were usually close to zero.

in the search bar will take visitor to a page that provides links to the product pages of products closely related to women's summer dress.

The recommendation system provides an additional option to visitors to find a desired product by navigating directly from a focal product's page to the displayed recommended products' pages (without having to go back to subcategory/category/home pages or type in the keyword in search bar). Thus, product recommendations reduce visitors' effort in finding the desired product on the website.

3.3 Product Recommendation Network

Showing related products on product pages creates a recommendation network of interconnected products on the retailer's website. An example of a part of this recommendation network is shown in Figure 2. The product page of the focal top in Figure 1 recommends four related tops, but the focal top may also appear as a recommended product in other products' pages. Accordingly, the focal top has outgoing recommendation links to its four recommended tops and it has several incoming recommendation links from other tops on whose page it appears as the recommended product as shown in Figure 2. Similarly, each of the four recommended tops have outgoing recommendation links to the products that are recommended on their pages and incoming recommendation links from the products on whose pages they appear as recommended products.

In a recommendation network, the connectedness of a product can be described by the number and strength of its recommendation links. The number of incoming recommendation links to a product, referred to as its indegree, indicates from how many other products' pages it can be reached on the website. A product with a large value of indegree would be displayed on a large number of other products' pages and thus would have a high likelihood of being viewed on the retailer's website. The number of outgoing recommendation links from a product, referred to as its outdegree, indicates the number of products that can be directly reached from the product's page. Over 97 percent of the total products on the retailer's website had outdegree of four.

As indicated in the previous section, the affinity score of a link in the recommendation network measures the strength of the relationship between the two products connected by the link. Thus, besides the numbers of incoming and outgoing links, the connectedness of a product in the recommendation network is additionally described by the average affinity scores of its incoming and outgoing links. Finally, all incoming links to a product are not similar, e.g., an incoming link from a popular product (a product with high indegree) may bring higher traffic to the product's page than an incoming link from a less popular product (a product with low indegree). To account for this fact, we also computed the PageRank of each product's page based on its interconnection with other products in the product recommendation network.[3] The PageRank of a product is an indicator of the traffic to its product page due to the network.[4]

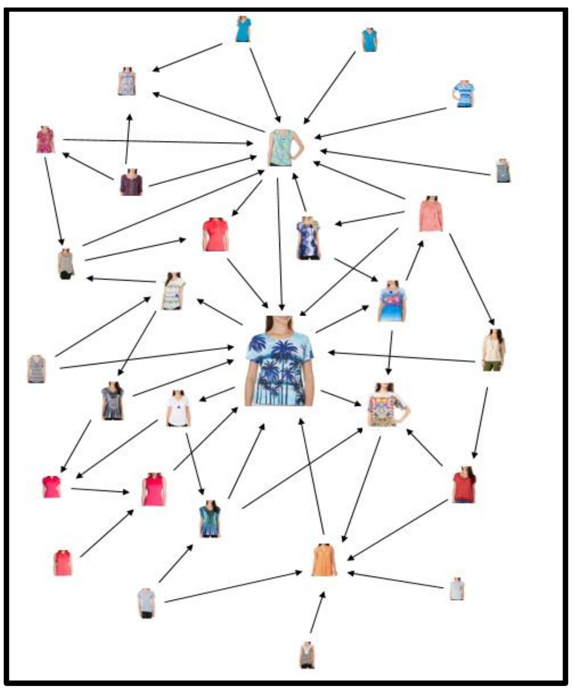

Figure 2: Product Recommendation Network

[3] PageRank is defined recursively, meaning that the PageRank of a node in a graph depends on the PageRank of all nodes that link to it (algorithm details available at https://www.mathworks.com/moler/exm/chapters/pagerank.pdf)

[4] We also compute the other forms of network centrality measures, such as closeness and betweenness centrality and conduct analysis on them. The details of this analysis are provided in Appendix C in this paper.

Between two products with the same PageRank, the product with the higher average incoming affinity score is likely to be viewed/bought more on the retailer's website. When a visitor views a close substitute recommended on a focal product's page (i.e., high affinity score between focal and recommended products), she may prefer to purchase the substitute instead of the focal product. Therefore, conditional on a focal product's page view, the probability of purchase of recommended products may increase with the average outgoing affinity score. Overall, while indegree (outdegree) of a product captures its number of connections, the average incoming (outgoing) affinity scores captures average strength of these connections in a recommendation network. The computation of network characteristics of the focal top in Figure 1 in the recommendation network is illustrated in Figure 3.

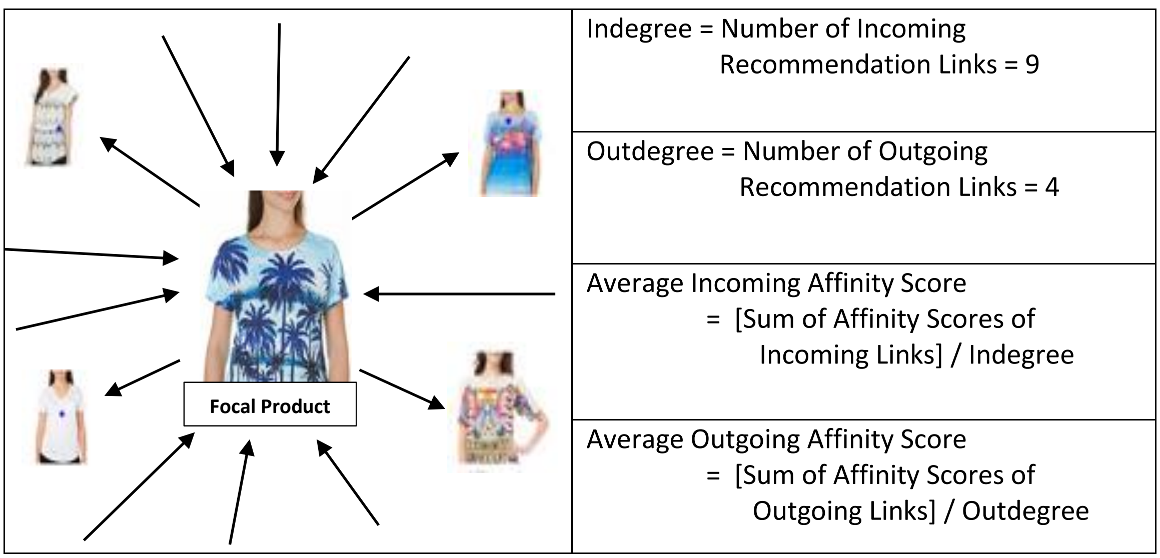

Figure 3: Network Characteristics of a Focal Product

3.4 Experimental Design

We wish to estimate the value of a product's recommendation links, defined as the total sales of the product and its recommended products attributable to their joint display on its product page. A visitor can directly visit and purchase a recommended product (e.g. discovering it via search or navigation options) instead of coming to its page from the focal product's page on the retailer's website. Therefore, to compute the value of a product's recommendation links, we only consider the sales of the product and its recommended products in the sessions in which the product's page is viewed. Moreover, if a visitor views the recommended product before viewing its focal product's page in a session, she does not discover the recommended product through its display on the focal product's page. Therefore, we only consider recommended product page views and sales in sessions in which they are viewed after viewing their focal products' pages.

Even if a visitor purchases the recommended product after viewing it on a product's page, such sales still cannot be definitively attributed to their display on the product's page for the following three reasons.

1. Many closely related products may appear on the same page under their product category or product subcategory main pages. For example, two closely related women's tops on Macy's.com appear on the women's top subcategory main page and one of these tops appears as the recommended product on other top's product page (see Figure A1 in Appendix A). Thus, a visitor may be already aware about the relationship between two women's tops after viewing them on the top subcategory main page.
2. Many keyword searches on the retailer's website may display closely related products on the search result page. For example, keyword search for "Women's tops" on Macy's.com results in the display of two closely related women's tops on the search result page and one of these tops also appears as the recommended product on the other top's product page (see Figure A2 in Appendix A). Thus, a visitor may be already aware about the relationship between two tops after viewing them on the search result page.
3. Many recommended products displayed on focal products' pages may be from the same designer with similar designs, colors, and styles that visitors may recognize even without their joint display. For example, on the product page of an Alfani designer women's top on Macy's.com, two other Alfani designer tops with similar styles are recommended (see

Figure A3 in Appendix A). Thus, visitors may recognize the relationship between similar products from the same designer independent of their display on each other's pages.

Therefore, to estimate the net effect on sales due to the visibility of recommended products on a product's page, the sales of recommended products due to their interrelationship with the product known to the visitors from above three reasons needs to be accounted for. We accomplish this by estimating the counterfactual sales of the product and its recommended products in presence of the above three reasons but in absence of their joint display on the product's page through an experiment.

Our experiment implements a standard A/B test in which the retailer created a control and a treated version of the product pages on its website. In the control version, no products were recommended on the focal product pages. In the treated version, up to four products with the highest affinity scores with the focal product and belonging to the same product subcategory were recommended on its page. An example of treated version of a product page is shown in Figure 1. The experiment was conducted on the retailer's website for nine weeks from 8th April 2015 to 9th June 2015. During this experiment, two-thirds of the total visitor sessions on the retailer's website were randomly assigned to either a control or a treated group.[5] If a session is assigned to the treated version, the visitor would remain in the treated group for all product page views in that session. Moreover, if a visitor comes to the retailer's website through the same machine on different days during the experiment period, her machine is recognized through cookies and she consistently remains in the same treatment group in all such sessions.[6]

4.0 Data Description

[5] The remaining one-third of the total visitor sessions were randomly left out of the experiment.
[6] The recommendation engine recognized the visitor machine through cookies and IP address. So if a visitor removes the cookie from her machine, s/he will not be recognized as the same visitor.

During the experiment period, there were a total of 1,307,191 visitor sessions with at least one product's page view. Out of these total sessions, 434,353 (33.30 percent) sessions were of control version, 435,411 (33.31 percent) sessions were of treated version, and the remaining 436,437 (33.39 percent) sessions were out of the experiment. Thus, we find that the two versions of recommendations were allocated to the visitor sessions with equal probability. However, we find that visitors view statistically higher number of web pages and product pages in the treated sessions as compared to the control sessions (see first two rows of the Table 1).

	Control sessions (N=434353)		Treated sessions (N=435411)		Diff in mean (*t-stats*)	
	Mean	St. Dev.	Mean	St. Dev.		
No. of web pages viewed per session	11.44	17.99	11.54	18.18	3.4	
No. of product pages viewed per session	2.77	4.21	3.10	4.91	38.6	
No. of sessions by a visitor	All sessions		Control sessions		Treated sessions	
	No. of sessions	% of sessions	No. of sessions	% of sessions	No. of sessions	% of sessions
1	536,609	61.7	267,545	61.6	269,064	61.8
2	112,957	12.99	56,291	12.96	56,666	13.01
3	57,077	6.56	28,841	6.64	28,236	6.48
4	36,743	4.22	18,360	4.23	18,383	4.22
>4	126,378	14.53	63,316	14.58	63,062	14.48
Total	869,764		434,353		435,411	

Table 1: Breakup summary statistics of visitor sessions

We further provide a breakup of total, control, and treated sessions based on the number of sessions by a visitor in Table 1. We find that 61.7 percent of sessions are by the visitors, who make only one visit to the retailer's website. However, we also find a significant number of sessions by the visitors who make more than four visits to the retailer's website (14.5 percent). Most importantly, Table 1 shows that the distribution of number of sessions by visitors are very similar across treated and control groups, which provides support to our randomization process.

During the experiment period, 37619 unique products were hosted on the retailer's website and out of those, product pages of 32173 products (85 percent of the products) were viewed in the visitor sessions.

4.1 Product Network Data

We collected data on the regular price, sale price, number of customer reviews, and mean review scores of the products on each day of the experiment period. For treated version of product pages of the products, we also collected information on the four recommended products, such as their ID, regular prices, sale prices, numbers of reviews, mean review scores, and the affinity scores with the product.

Mean values for 37619 Products	Percentile values						
	0	25	50	75	90	95	100
Regular price (US $)	1.99	20	36	54.99	79.5	118	680
% Price discount	0	9.1	25.2	34.6	40.0	44.8	82.9
No. of reviews	0	0	0	0	0.8	1.7	473.5
Review score	0	0	0	0	3.2	4.1	5

Table 2: Product characteristics

Since the price, number of reviews, and review score for products may change over the experiment period, we computed the mean values of these variables over the experiment period for all products and report their distribution at different percentiles in Table 2. The main takeaways from Table 2 are that: products are diverse in terms of their prices; are often discounted (median discount 25%); and no customer reviews are available on most product pages.

As illustrated in section 3.2, recommending related products in the treated version of product page results in the creation of a product recommendation network on the retailer's website. For this product network, we collected data on the network characteristics (indegree, PageRank, outdegree, and affinity scores for the incoming and outgoing links) on each day during the experiment period. The affinity scores between focal-recommended product pairs in the networks may change from one day to another based on how frequently different products are viewed, carted, abandoned, and bought with other products. Therefore, the network characteristics of the products would change over time. To get a sense of how much a product's network characteristics differ across our sample of products and within products over the period of the experiment, we computed their mean values and the range of their variations and report their distributions in Table 3.

37619 products	Percentile values						
	0	25	50	75	90	95	100
Mean indegree	0	0	1.1	3.7	8.1	13	299.5
Range of indegree	0	0	4	8	16	24	452
Mean PageRank	4.6E-06	4.9E-06	7.6E-06	2.7E-05	7.3E-05	1.2E-04	1.4E-03
Range of PageRank	0	5.1E-07	1.5E-05	7.3E-05	1.7E-04	2.6E-04	2.6E-03
Mean avg. incoming affinity score	0	0	11.0	40.3	105.8	171.8	9999999
Range of avg. incoming affinity score	0	0	33	87.9	193.6	291	9999999
Mean outdegree	1	4	4	4	4	4	4
Range of outdegree	0	0	0	0	0	0	3
Mean avg. outgoing affinity score	0	4.1	23.3	76.8	203.1	372.8	9999998
Range of avg. outgoing affinity score	0	11.7	36.7	94.5	209.4	331.0	5000040

Table 3: Summary statistics of network characteristics

We observe high variations in both the mean values of indegree (PageRank) across products and indegree (PageRank) within products over the experiment period. About half of the products have a mean value of indegree of one or less in the network and their indegree values change by up to 4 during the experiment period. The outdegree of about 97 percent of the products was four, as most of the products' pages on the retailer's website recommend exactly four related products, and the outdegree value did not change over the experiment period. We also observe large variations in both, the mean incoming (outgoing) affinity scores across products and within products over the experiment period. Overall, we find significant variations in network characteristics across products at a time and within products over time.

4.2 Visitor Session Data

Visitors viewed the product pages of 32173 unique products (out of the total 37619 products) in 869,764 sessions (434,353 control sessions and 435,411 treated sessions) on the retailer's website during the experiment period. In order to estimate the value of a product's recommendation links, we break down a visitor session into a sequence of product page views. As a result, we collected data on 2,326,402 product page views in 869,764 visitor sessions. Out of these, 1.086,222 (47 percent) product page views were in control sessions and 1,240,180 (53 percent) product page views were in treated sessions. Thus, we observe a greater number of product page views per session in the treated version. This suggests that the increased visibility of relevant products in the treated version reduces search costs and drives greater product exploration.

	Mean	St. Dev.	Mean	St. Dev.	Diff. in means
	Control sessions		Treated sessions		*t-value*
No. of product page views	1086222		1240180		
No. of recommended product page views after product's page view	0.176	0.381	0.267	0.442	167.58
No. of products purchased	0.064	0.246	0.058	0.237	-16.64
$ value of products purchased	1.825	8.831	1.667	8.490	-13.90
No. of recommended products purchased	0.018	0.131	0.020	0.140	13.65
$ value of recommended products purchased	0.482	4.484	0.565	4.899	13.50
Probability of only product purchase	0.056	0.230	0.051	0.220	-17.28
Probability of only recommended product purchase	0.011	0.102	0.013	0.114	18.11
Probability of both product purchase	0.007	0.083	0.007	0.083	-1.23

Table 4: Summary statistics of visitor session data

For each product's page view in visitor sessions, we collected information on number of page views of the recommended products after the product's page view in the session, number ($ value) of products purchased, and number ($ value) of recommended products purchased. [7] We

[7] It is possible that some pairs of products may have reciprocal focal-recommended product relationship, e.g., a product B may appear as a recommended product on a product A's page and A may, in turn, appear as a recommended product on B's page. In such cases, we record the product page views in a session as follows.

- If a visitor first views A's page and then views B's page in a session: For such a session, we record A's page view (with recommended B's page view) and once as B's page view (with no recommended product's page view).
- If a visitor views product pages of A and B several times in a session, such as A-B-A: Such a session will be counted once as A's page view (with recommended B's page view), once as B's page view (with recommended A's page view), and once as only A's page view.

report the summary statistics of these variables for the product page views in treated and control sessions separately and also report the *t-value* for the difference in their mean values in Table 4.

From Table 4, we make several interesting observations. First, we find a statistically greater number of recommended product page views after the product's page views in the treated sessions than that in the control sessions (0.267 versus 0.176). We further find that conditional on viewing the product pages, a statistically smaller number ($ value) of products and a statistically greater number ($ value) of their recommended products are purchased in the treated sessions as compared to that in the control sessions. Interestingly, we find on average 0.176 page views (67 percent of 0.267 page views in treated session) and $0.482 worth of purchases (86 percent of $0.562 purchases in treated session) of recommended products after the product's page view in control sessions in which the recommended products are not displayed on the product's page. This indicates that visitors recognize the relationship between the products and their recommended products due to the reasons enumerated in section 3.4. Therefore, not accounting for this fact could lead to overestimation of the value of recommendation links.

Visitors may purchase the recommended products with the product if they perceive them as similar and purchase recommended products in place of the product if they perceive them as substitutes. We examine this issue in Table 4 by looking at how the probabilities of purchasing only product, only recommended product, and both products vary across treated and control sessions. We find that the probability of purchasing both products in the treated and control sessions are statistically similar, but a statistically lower (higher) probability of purchasing only product (recommended product) in treated session as compared to that in control sessions. This shows that on average visitors view products recommended via the current recommendation system as substitutes to the products. This is consistent with the design of the recommendation system.

5.0 Empirical Analysis and Results

In this section, we compute the value of a product's recommendation links and examine how it varies with its network characteristics. We first compute the value of a link in a model-free manner, and then present a robust fixed-effects regression specification for the same.

5.1 Value of a Product's Recommendation Links

When we are measuring the value of recommendation link, it is important to ask "compared to what?" We could compare the value of showing a recommendation link with that from showing a bestseller, a new product in the catalog, a random product from the catalog or showing no product on the focal product's page. Showing a bestseller, a new product, or a random product on a focal product's page is rarely meaningful in practice; showing such products on the homepage often makes sense and offers value but showing them on a specific product page provides little value to the consumer. Hence, it is a practice that is not used by most retailers. Accordingly, we focus on the value of a recommendation link against the alternative of not showing any other product on the product's page. This is in fact the control group in our study and it assesses the direct value of recommenders against the dominant alternative prior to their adoption by retailers. One limitation with this approach is that it doesn't separately quantify the added value of recommending closely related products versus the value of merely exposing the consumer to *any* product in the retailer's catalog. We address this question in detail in Section 5.4.

Given the above, the value of a product's recommendation links is the daily $ sales of a product and its recommended products on the retailer's website attributable to their joint display on the product's page, and can be computed as follows:

$$\begin{aligned}&[\text{Number of daily page views of the product * Mean \$ sales of the product}\\&\quad \text{and its recommended products} \mid \text{product's page view}]_{\textit{Treated Sessions}}\\&-[\text{Number of daily page views of the product * Mean \$ sales of the}\\&\quad \text{product and its recommended products} \mid \text{product's page view}]_{\textit{Control Sessions}} \qquad \text{---- (1)}\end{aligned}$$

5.2 Model-free Evidence of Value of a Product's Recommendation Links

In specification (1), we already know the mean $ sales of a typical product and its recommended products conditional on its page views in the treated and control sessions (computed from the visitor session data in Table 4). But, we do not know daily average number of page views of products, which we compute by aggregating the daily page views for each product for the treated and control sessions separately. If a product is viewed in only one type of session on a day, we put the number of page views of that product in the other type of sessions on that day as zero. We finally obtained a total of 3,958,166 product-day observations, 1,979,803 observations each for the two types of sessions in which 37619 products' pages were viewed over the 63 days of experiment period. We find that a product's page was viewed on average 0.714 and 0.62 times per day, respectively, in the treated and control sessions (the corresponding standard deviation values were 4.32 and 4.03, respectively). Thus, we find a higher average number of daily page views of products in treated sessions.

	Control	Treated	Diff.
Daily avg. number of product page views	0.620	0.714	0.094
Mean $ sales of product |product's page view	1.825	1.667	-0.158
Mean $ sales of recommended products | product's page view	0.482	0.565	0.083
Daily avg. $ sales of product	1.132	1.190	0.059
Daily avg. $ sales of recommended products	0.299	0.404	0.105
Daily avg. $ sales of product + recommended products	1.431	1.594	0.163

Table 5: Daily $ value of a product's recommendation links

Next, we use specification (1) to compute the daily $ sales of an average product and its recommended products attributable to their joint display on the product's page and report it in Table 5. We find a higher average $ value of recommended products sales ($0.565 versus $0.482) and a lower average $ value of the product's sales ($ 1.667 versus $1.825) given product's page view in the treated sessions as compared to that in control sessions. This suggests that visibility of recommended products on a product's page results in visitors substituting their purchases of the product with recommended products. However, we also find a greater average number of product page views in treated sessions as compared to that in control sessions (0.714 versus 0.62), which mitigates the effect of substitution. Overall, we find gains of $0.059 and $ 0.105, respectively, in the daily sales of a product and its recommended products in treated sessions as compared to that in control sessions. This translates into a net sales gain of $ 0.163 due to the visibility of a product's recommendation links on the retailer's website, which is 11.4 percent increase over the average sales of $ 1.431 in control sessions. Thus, the daily average value of recommendation links in our setup is $ 0.163 per product. However, as we highlight below, this gain does not accrue equally among all products.

	Products in bottom quartile (Mean PageRank < 0.0000049)			Products in top quartile (Mean PageRank > 0.000027)		
	Control	Treated	Diff.	Control	Treated	Diff.
Daily avg. number of product page views	0.181	0.184	0.003	1.418	1.755	0.337
Avg. $ sales of product | product's page view	1.285	1.220	-0.066	2.066	1.814	-0.252
Avg. $ sales of recommended products | product's page view	0.408	0.415	0.007	0.518	0.629	0.110
Daily avg. $ sales of product	0.233	0.224	-0.008	2.929	3.183	0.254
Daily avg. $ sales of recommended products	0.074	0.076	0.003	0.735	1.103	0.368
Daily avg. $ sales of product + recommended products	0.307	0.301	-0.006	3.664	4.286	0.622

Table 6: Variations in the value of a product's recommendation links with PageRank

We further examine how the value of a product's recommendation links vary with its network characteristics. For this analysis, we divide the 37619 products into the top and bottom quartiles based on their average PageRank during the 63 days of experiment period. The average PageRank values for products in the bottom and top quartiles were, respectively, less than 0.0000049 and more than 0.000027 (see Table 3). We computed the total sales for the product and its recommended products for products in the two quartiles on similar lines as in Table 5 and report it in Table 6. We find a gain of $ 0.622 from recommendation links of products in the top quartile but almost no gain in sales from recommendation links of products in the bottom quartile. Thus, products with high PageRank benefit from recommendations but not the products with low PageRank.

5.3 Econometric Analysis

In the previous section, we computed the value of a product's recommendation links by simply taking the difference of mean sales of the product and its recommended products in treated and control sessions. However, a product's sales on the retailer's website is affected by a variety of product- and time-specific factors. Thus, in the following we use regression specifications to precisely estimate the value of a product's recommendation links on the retailer's website.

5.3.1 Regression-based Evaluation of Value of a Product's Recommendation Links

To estimate the value of a product's recommendation links, we use the following specifications:

$$LNPViewpday_{ikt} = \beta_1 \times Treat_{ikt} + \alpha_{it} + \varepsilon_{ikt} \quad ; \qquad \text{-- (2)}$$

$$LPSales\ (LRPSales) |\ PView_{ijt} = \beta_1 \times Treat_{ijt} + \alpha_{it} + \varepsilon_{ijt} \quad ; \qquad \text{-- (3)}$$

where, *i* denotes the 37619 products, *t* denotes 63 days of experiment, *j* denotes page views of these products in 869,764 visitor sessions, and k denotes the session type (treated or control). In specification (2): $LNPViewpday_{ikt}$ denotes the log of number of page views of product *i* in sessions of type *k* on day *t*; $Treat_{ikt}$ is an indicator variable equal to one if session type *k* with product *i*'s page view on day *t* is treated and zero otherwise. In specification (3): $LPSales\ (LRPSales)\ |\ PView_{ijt}$ denotes the log of $ value of product's *i* sales (recommended products' sales if viewed after the product's *i* page view) given product's *i* page view *j* in a session on day *t*;[8] $Treat_{ijt}$ is an indicator variable equal to one if product's *i* page view *j* on day *t* is in a treated session and zero otherwise. Parameter α_{it} denotes product-day fixed effects. In specifications (2) and (3), product-day fixed effects account for the unobserved product-day level factors that may affect the sales of products such as any systematic promotions offered by the retailer based on the sales of the product at a

[8] We take the log of dependent variables in specification (2) and (3) on account of their highly skewed distribution as shown in Table 3. We obtain qualitatively similar estimates without log transformation of dependent variables.

specific time. For example, orange color products may be offered preferential placement or price discounts on the retailer's website around Halloween time or a celebrity may promote a dress at a time when it is not selling well.

	Log of number of sessions with product view per day		Log of Product sales | product view		Log of Recommended product sales | product view	
	Coeff. Est.	St Err.	Coeff. Est.	St Err.	Coeff. Est.	St Err.
Treat	0.075^{***}	0.001	-0.019^{***}	0.002	0.09^{***}	0.001
Product-day fixed effect	Yes		Yes		Yes	
N (No of Product-days)	3958166 (1979083)		2326402 (701730)		2326402 (701730)	

∗∗∗, ∗∗, and ∗ denote statistically significant at α=0.01, 0.05, and 0.10 levels (two-sided test), respectively
Standard errors cluster corrected at product-day level

Table 7: Regressions estimates for the value of a product's recommendation links

The coefficient estimates of specification (2) & (3) are reported in Table 7. We find a positive and significant coefficient value of 0.075 for *Treat* variable in specification (2), which indicates a 7.5 percent higher page views of the products in treated as compared to that in control sessions. We find a negative and significant coefficient value of 0.019 for *Treat* variable for the product sales and a positive and significant coefficient value of 0.09 for its recommended product sales in specification (3). This indicates that conditional on the product's page view, there is a 1.9 percent lower sales for the product and a nine percent higher sales for the recommended products in treated sessions as compared to that in control sessions.

It is possible that in treated sessions with the products' page views, visitors may purchase the products and/or their recommended products instead of purchasing *other* products offered on the retailer's website. If this is the case, then the loss in sales of other products subtracted from the gains in sales of the products and their recommended products in the treated sessions would be the true value of a product's recommendation links. To check for this possibility, we estimate specification (3) with $ value of other products' sales in a session as the left hand side variable. We find a negative and insignificant coefficient value of 0.05 (*t-value* =0.62), which indicates statistically similar sales of other products in the two types of sessions. Thus, specification (1) computes the true value of a product's recommendation links in our setup.

We further used the estimated coefficients in specifications (2) and (3) to estimate the net impact of recommendations on the retailer's aggregate sales on the website. On average, we find an additional \$0.03 product sales and \$ 0.085 recommended products sales in treated sessions as compared to the control sessions per day, which translates into an additional \$ 0.115 sales per day in the treated sessions. This value is close to \$ 0.16, the value we obtained in Table 5 with the difference of mean daily sales of these products in treated and control sessions. As only one-third of the total sessions were treated with recommendations in the experiment, the total gains in sales with treatment of all session would be \$ 0.115*3 = \$ 0.345.[9] For approximately 35000 products hosted on the retailer's website per day, this translates into a total sales gain of \$ 0.345*35000 = \$ 12075 per day or \$ 4.4 million per year. Thus, 11 percent of the retailer's annual online sales of over \$ 40 million can be attributed to the recommendations of related products on the product pages.

5.3.2. Variation in Value of a Product's Recommendation Links with Network Characteristics

To understand the effect of network characteristics of a product on its recommendation links, we analyze how these characteristics could affect: (1) the daily number of its page views and (2) its sales (and its recommended products sales) conditional on its page view. The daily number of page views of a product would be influenced by its PageRank and its average incoming affinity score. A product with high PageRank would be accessible from a large number of other popular products and thus its product page has a higher likelihood of being viewed. Similarly, a product with high average incoming affinity scores has a higher likelihood of being visited from its related products' pages due to their close relationship. Once a visitor comes to a product's page, her propensity to

[9] Note from section 3.3 that only one-third of the total session were allocated to the treated version in the experiment.

purchase the product and its recommended products would be influenced by its outdegree and average outgoing affinity score. Recommending close substitutes, having a high average outgoing affinity score with the product, on a product's page may increase the purchase of recommended substitutes at the expense of its purchase.

Accordingly, we examine the effect of a product's network characteristics on the value of its recommendation links with the following specifications:

$LNPViewpday_{ikt} = \beta_1 \times Treat_{ikt} + \beta_2 \times\ (Treat.Pgrank)_{ikt} + \beta_3 \times (Treat.LInAffinity)_{ikt} + \alpha_{it} +\ \varepsilon_{ikt}$; --- (4)

$LPSales\ (LRPSales)\ |\ PView_{ijt} = \beta_1 \times Treat_{ijt} + \beta_2 \times (Treat.LOutAffinity)_{ijt} +\ \ \alpha_{it} + \varepsilon_{ijt}$; --- (5)

where, the subscripts *i*, *j*, *k*, and *t* have the same meaning as in specifications (2) and (3). Variables $Pgrank_{it}$, and $LIn(Out)Affinity_{it}$, respectively, denote PageRank, and log of average incoming (outgoing) affinity scores of product *i* on day *t*.[10] Variables *Treat.Pgrank*, and *Treat.LIn(Out)Affinity* are, respectively, the interaction terms for these variables with *Treat*. The remaining variables have the same meaning as in specifications (2) and (3). Variable *Treat.Outdegree* is not separately included in specifications (5) because it is highly collinear with *Treat*.[11]

The retailer may systematically promote high selling products that may also have high network characteristics (such as PageRank) on its website. Even though, we did not find any correlation in price promotion of a product with its network characteristics in our setup, non-price promotions on a product may still be correlated to its network characteristics.[12] For example, the

[10] The logarithm of Average incoming affinity scores and Average outgoing affinity scores variables are taken due to their highly skewed distribution.
[11] Note that 97 percent of the product have *outdegree* of 4 that remains constant during the experiment period.
[12] If the retailer offers price promotions selectively based on the lower (higher) sales of a product due to its poor (superior) network characteristics, then there should be a negative (positive) correlation between the price discount and network characteristics of products. To check this in our data, we examine the correlation between the price discount offered for a product on a given day with its centrality measures on the same day. We find almost no correlation between price discount and all network centrality measures in our data. Specifically, the correlation coefficient for discount with Indegree, PageRank, Closeness centrality and weighted betweenness centrality are (0.0055), (0.0069), (0.011) and (0.013) respectively.

retailer may offer preferential placements to a high selling product on its website. Such systematic promotions to products with high network characteristics at a time are product-day level unobserved factors and are accounted for by the product-day fixed effects in specifications (4) and (5).

	Log of daily number of sessions with product's page view	Log of Product sales | product's page view	Log of Recommended product sales | product's page view
Treat	-0.008*** (0.0004)	-0.002 (0.011)	-0.002 (0.001)
*Treat * PageRank*	346.53*** (10.31)		
*Treat * Log avg. incoming affinity score*	0.011*** (0.0003)		
*Treat*Log avg. outgoing affinity score*		-0.003** (0.001)	0.012** (0.0004)
Product-days fixed effects	Yes	Yes	Yes
N(No. of product-days)	3958166 (1979083)	2326402 (701730)	2326402 (701730)

∗∗∗, ∗∗, and ∗ denote statistically significant at α=0.01, 0.05, and 0.10 levels (two-sided test), respectively
Standard errors cluster corrected at product day level are in parentheses

Table 8: Regression estimates for network characteristics

In Table 8, we find a negative and significant coefficient value of 0.008 for *Treat* indicator variable in specification (4). This coefficient value means that products with PageRank=0 and no incoming links (i.e. Incoming affinity score=0) are on average viewed in 0.8 percent fewer times per day in treated sessions as compared to that in control sessions. We further find a positive and significant coefficient values 346.5 for the interaction terms of *Treat* with *PageRank*, which translates into an additional 0.1 percent page views in treated sessions with 10 percent increase in PageRank value over the average value of 0.000029 in our sample. Similarly, we find a positive and significant coefficient value of 0.011 for the interaction term of *Treat* with log of average incoming affinity scores, which translates into an additional 0.1 percent page views with 10 percent increase in average incoming affinity score. Thus, we find that average incoming affinity score (strength of links) and the PageRank (location of links) of a product have similar magnitude of impact on its number page views. Overall, these results suggest that additional visibility of a product on other products' pages and on the pages of strongly related products increase the likelihood of its page view.

For specification (5), we find that the coefficient of interaction term of *Treat* with log of average outgoing affinity scores has a negative and significant coefficient value of -0.003 for product sales and a positive and significant coefficient value of 0.012 for its recommended products' sales, which suggests that showing close substitutes on a product's page would increase the sales of recommended substitute at the expense of its sales.

Our econometric analysis confirms two main findings. First, recommending substitute products indeed cannibalizes sales of focal products in favor of the recommended products but such cannibalization is more than offset by the increase in number of page views of focal products. The net impact is an increase in total sales of the focal and its recommended products, on average. Second, the strength of a focal product's incoming links (average incoming affinity score) is one of the main drivers for increase in its page views, and hence a major determinant of the value of recommendation links in a product network.

5.4 Recommendation Versus Exposure Effect of Recommendation Systems

We observe that recommendation links on a focal product page drive up the sales of recommended products. The mechanism question that we examine here is twofold. First, the estimated value of a product's recommendation links could come from two factors: (1) effect of merely showing products on a focal product's page (exposure effect) and (2) effect of showing closely related products on a focal product's page (recommendation effect). The relative strengths of these two effects, and therefore the treatment effect of recommendation system, may vary based on the characteristics of the recommended products. In short, the first mechanism question is whether the increase in sales is due to a mere exposure effect or is there a recommendation effect as well. In addition, because recommendations are based on past views and sales, popular items may end up being recommended more. So, the second mechanism question is whether the effects differ in showing popular vs unpopular product. Finally, these two mechanisms may even interact. That is, the exposure effect is likely to be higher if popular (presumably high quality) products are recommended on a focal product's page and the recommendation effect is likely to be higher if strongly related products are recommended on a focal product's page.

The strength of relationship between a focal and its recommended products is measured by their affinity scores, a high average affinity score indicates a strong relationship between products and vice versa. In Table 3, we have already shown the distribution of mean values of average outgoing affinity scores for products during our experiment. We utilize the quartile split of average outgoing affinity scores of products to categorize very strong, strong, weak, and very weak relationship with their recommended products. Next, we computed the total daily sales of recommended products for each product on the retailer's website on each day of the experiment. We use the quartile split of the total daily sales of recommended products to categorize them into very popular, popular, unpopular, and very unpopular recommended products. In Table 9, we report the cutoff values for these categories.

Popularity categories	Average daily sales of recommended products (numbers)	Strength of recommendations Categories	Average outgoing affinity scores
Very Unpopular	0	Very Weak	<=4.1
Unpopular	0< and <=2	Weak	4.1< and <=23.3
Popular	2< and <=5	Strong	23.3< and <=76.8
Very Popular	>5	Very Strong	>76.8

Table 9: Product categories based on popularity and strength of recommendations

Since the daily sales for very unpopular recommended products are zero in all sessions, we cannot estimate the effect of recommendations for this category of products. So we clubbed the unpopular and very unpopular categories together in one unpopular recommended product category. Accordingly, we created six samples of focal-recommended product pairs by combining three popularity categories (very popular, popular and unpopular) and two strength of relationship categories (very weak and very strong). Comparing the sales of recommended products for these samples in the treated sessions as compared to that in control sessions would reveal the relative strength of recommendation and exposure effects.

We estimate specification (3) for recommended product sales for the six samples of treated and control sessions data. Note that in each of such samples, both treated and control sessions are included in which the focal products with a specific category of recommended products were viewed, and thus comparison of recommended product sales between the two types of sessions would indicate the value of recommending that specific category of products versus no recommendations. The coefficient estimates for treatment effect for sales of recommended products from specification (3) in each of these samples are reported in Table 10.

We find statistically higher sales in treated sessions for strongly related recommended products of all popularity categories (see row two of Table 10), but not for very weakly related recommended products in any popularity category (see row three of Table 10). To test whether insignificant estimates for very weakly related recommended products are due to small sample sizes, we additionally estimate specification (3) on the combined samples of weakly and very weakly related recommended products and find similar insignificant estimates (see row four of Table 10).

Treatment effect of Recommendations	Very popular	Popular	Unpopular
Very Strong	0.012^{***} (0.002) [425893 (48703)]	0.0098^{***} (0.002) [407783 (77747)]	0.003^{***} (0.001) [889481 (260833)]
Very Weak	0.006 (0.015) [3163 (1785)]	-0.007 (0.027) [3973 (2457)]	-0.005 (0.005) [41091 (27865)]
Weak + Very Weak	0.009 (0.011) [9495 (4451)]	0.016 (0.012) [14700 (7674)]	-0.0006 (0.003) [161317 (103744)]

$***$, $**$, and $*$ denote statistically significant at α=0.01, 0.05, and 0.10 levels (two-sided test), respectively. Coefficient estimates (standard errors cluster corrected at product-day level) [No. of observations (No. of product-days)]

Table 10: Treatment effect of recommendations for different product categories

Overall, recommending weakly related products does not drive sales of the recommended products. This suggests that the effect of recommendation system in our setup is primarily due to the recommendation effect. Further, recommending popular products is not helpful unless the recommendations are relevant, which suggests absence of exposure effect. So, popularity of recommended products is by itself not valuable. But if recommended products are both popular and relevant, that appears to deliver the greatest increase in sales of recommended products.

5.5 Robustness checks

We further conducted several robustness checks of our central results, which we describe below.

1. Our main analysis focused on two main network centrality measures, namely a product's degree and PageRank. In Appendix B, we check whether our results are robust to the other types of network centrality measures. We find that different types of centrality measures individually yield qualitatively similar results. We further find that our results are robust to inclusion of a combination of uncorrelated centrality measures (PageRank and Weighted betweenness).
2. If a pair of selected products (say A and B) with reciprocal focal-recommended product relationship (B is recommended on A's page and vice versa) are seen several times in a session (for example A-B-A), our specification may simultaneously pick up the cannibalization of A's sales at the expense of B's sales and of B's sales at the expense of A's sales. We verify in Appendix C that our results are robust to the exclusion of sessions with page views of products with reciprocal relationships.
3. Our data are derived from a website that sells wide range of products, from apparel and related fashion products to luggage and home furnishings. We check whether our results are robust across fashion products (more hedonic in nature) and non-fashion products (more utilitarian in nature). In Appendix D, we classify different product categories into fashion and non-fashion products, and estimate our specifications on these categories separately. We find that our main results hold for both categories of products. This, suggests that our results on economic value of recommendation links are generalizable, even though the magnitudes of results may vary across different product categories.

6.0 Managerial Implications: Analysis of Sponsored Recommendations

Our study reveals the potential value of recommendation links for product success. A natural question is how can producers and retailers think strategically about our findings and what kinds of actions can managers take? One relevant action can be seen in the context of sponsored recommendations. Amazon manages an advertising platform to allow advertisers to pay Amazon to place their products on the product pages of other relevant products. At the time of writing, most product pages on Amazon.com feature "sponsored products related to this item." Similarly, Alibaba also features sponsored recommendations on product pages on its website. For retailers, this allows them to monetize product pages better and offer relevant alternatives to consumers beyond the three or four products selected by the recommender algorithm (organic recommendation). For a producer, the value of the advertising platform is that they can advertise alongside relevant products. If their product is a relevant substitute or complement to a focal product (e.g. have a high affinity score with the focal product), advertising it on the focal product's page can help increase exposure.

The question, however, is what kinds of information and tools should retailers provide to producers so they can make good advertising decisions (i.e., identify focal products' pages on which to advertise). In sponsored search, companies like Google offer many tools to advertisers to help them identify search keywords and bid strategies. Similar decision support tools for advertisers are likely to be highly valuable in online retailing as well. Specifically, how can a producer estimate the value it is likely to derive from an additional inlink to its product from sponsored display on a focal product's page (which in turn is helpful in determining bids in ad auctions)? In this section, we illustrate how our analysis provides insights in this regard. Our discussion is motivated by the observation that many products that are not displayed on a focal product's page actually have a high affinity score with the focal product but not high enough to make it as one of the top four recommendations.

6.1 Retailer's Problem in Sponsored Recommendations

The retailer would like to alert producers about opportunities to advertise their products as sponsored recommendations. To do so, it needs to identify the best candidates and help them quantify the potential value from advertising. Consider a focal product page. From the retailer's perspective, it's worth asking which products are good candidates for sponsored display on that page. The retailer can alert producers of those products of the opportunity through its ad platform. One way to identify such products is to find all products with high affinity with the focal product that are not already recommended on the focal product's page. We illustrate this in Table 11 for a sample focal product (id= 320327).

Product id	Affinity score with focal product	Rank (by affinity score)	Original PageRank	Original Avg. In. Affinity score	New PageRank
346093	3690.32	5	1.65E-04	172.9	4.90E-04
291779	1953.7	6	1.83E-04	38.4	3.33E-04
338363	1936.17	7	8.50E-06	20.3	8.53E-06
343703	1530.63	8	1.36E-05	66.2	6.55E-04
333511	1018.1	9	9.68E-05	34.6	9.68E-05
333484	812.03	10	6.10E-06	19.7	1.73E-05
346085	523.28	11	2.13E-05	26.4	2.71E-04
340484	440.06	12	9.90E-06	6.1	1.99E-04
282713	436.59	13	5.29E-05	89.4	1.40E-04
343677	373.57	14	6.90E-06	22.1	3.69E-04
252950	322.19	15	4.53E-05	55.8	1.43E-04
Product id	**New Avg. In. Affinity score**	**Increase in No. of sessions**	**Purchase probability**	**Product price**	**Revenue gain**
346093	492.75	0.35	0.03	70	0.72
291779	55.52	0.16	0.19	65	2.03
338363	658.91	0.06	0.04	65	0.17
343703	432.31	0.69	0.11	70	5.33
333511	110.24	0.02	0.08	65	0.11
333484	415.86	0.07	0.05	65	0.23
346085	150.58	0.29	0.06	65	1.04
340484	92.87	0.24	0.10	65	1.54
282713	116.09	0.09	0.09	65	0.54
343677	139.23	0.40	0.15	65	3.92
252950	82.45	0.11	0.04	65	0.28

Table 11: Revenue gains from sponsored display of products on product (id =320327) page

The eleven rows in the table correspond to eleven products that are ranked from 5 to 15 in terms of their affinity scores with the focal product (products ranked in the top 4 are already shown as organic recommendations and hence not considered here). The producers of these eleven products are good candidates as advertisers.[13] Using our model, we can estimate (approximately) the likely value to an advertiser if it wins the auction and its product is displayed as the fifth recommended product, albeit as a sponsored recommendation. Below we show this estimation process in steps

- If product id = 343703 is advertised on the product page of product 320327, it now receives a new inlink. This additional inlink will change its PageRank and average incoming affinity score, and likely increase the number of sessions with its page view. We use specification (4) to estimate the number of daily sessions in which this product will be viewed based on its new PageRank and average incoming affinity score. Table 11 reports the original and new values of PageRank and average incoming affinity scores for product 343073 and reports the change in number of daily sessions ("*Increase in Sessions*") after its sponsored display.
- The net change in sales (revenue gain) for product 343703 is computed by multiplying the increase in its number of daily sessions with the purchase probability conditional on its page view and its purchase price. We use the average purchase probability of product 343073 in last seven days for this computation. Product 343073 gains an additional daily revenue of $5.33 from its sponsored display on focal product 320327 page.
- Likewise, the additional revenue for other ten products from their sponsored display at the product page of product 320327 are computed and listed in Table 11.

[13] There is no reason why fifteen is a good cutoff for this discussion. We use the top fifteen for illustration because we had the affinity scores of up to fifteen related products for focal products in our present setup. However, an advertiser can easily be more stringent or liberal in its selection criterion.

This way the retailer can shortlist candidate advertisers for all products and signal to them the potential value from advertising.[14]

6.2 Advertiser's Problem in Sponsored Recommendations

Advertisers need to look at all potential sponsored recommendations opportunities collectively and identify which ones are most appealing. From Table 11, we find that among the candidate products for sponsored display on product 320327 page, product 343073 gains the maximum additional revenue. However, from the perspective of the producer of product 343703, there may be multiple products' pages where the producer could advertise its product. Identifying the list of all such candidate products and the potential value of advertising product 343703 on those product pages is the relevant managerial question for a producer.

Focal product id	Affinity score with advertiser	Rank (by affinity score)	New PageRank of advertiser	New Avg. In Affinity score of advertiser	Increase in No. of sessions	Revenue gain
320291	2630.36	6	0.000738	707.24	0.79	6.06
320327	1530.63	8	0.000655	432.31	0.69	5.33
319407	864.58	8	0.00039	265.80	0.41	3.17
320313	1397.62	7	0.000303	399.06	0.33	2.54
255159	982.32	6	0.000194	295.23	0.21	1.63
339019	1864.73	7	0.000132	515.84	0.16	1.22
323074	1127.52	7	0.000131	331.53	0.15	1.15
351476	1082	7	0.000115	320.15	0.13	1.02
344055	700.21	9	7.12E-05	224.71	0.08	0.62
346093	2681.06	5	4.39E-05	719.92	0.07	0.57

Product 343073 has the original PageRank= 0.0000136; original Avg. InAffinity score =66.21; price = $70; and purchase probability conditional on its page view in last seven days =0.11

Table 12: Revenue gains from advertising product 343703 on other product pages

[14] Our analysis assumes that specification (4), which was estimated based on data from organic recommendations on the retailer's website, can be used to estimate the impact of a sponsored recommendation. This is ultimately a simplification and therefore the values in Table 9 & 10 are approximations at best. Note that there is no reason why an advertiser has to make this assumption in practice. Once the ad platform has been running for some time, the retailer can estimate a model that computes the number of sessions for a product as a function of the PageRank and InAffinity associated with organic recommendation links *as well as* the PageRank and InAffinity associated with sponsored links. The simplifying assumption has been made here because there were no sponsored recommendations on our partner website. So our analysis is illustrative at best and meant to mainly demonstrate a framework that retailers and advertisers can use in sponsored recommendation markets.

We begin by identifying all product pages for which product 343073 is ranked between 5 and 15 in terms of affinity score. These are all good target pages where the producer might consider advertising product 343703. In Table 12, we show the top ten options for advertising product 343073.[15] Each row in the table corresponds to a product ("focal product") on whose page product 343073 can be advertised as a sponsored recommendation. Potential revenue that product 343073 gains from advertising on different products' pages is calculated using the steps described in Table 11. Note that the producer of product 343703 can obtain the highest revenue gain of $ 6.06 per day by advertising on the page of product 320291. Also notice that there is considerable heterogeneity in the potential revenue gains even though product 343703 has high affinity with all of the target products listed in Table 12. This heterogeneity arises because some of the target pages have very high PageRank values, and recommending product 343073 on such products' pages can substantially increase the exposure and traffic for the advertised product. Such information can be very helpful to advertisers in developing an efficient advertising strategy i.e. identifying target pages on which to advertise and allocating the ad budget among them.

We acknowledge here that the users may accord lower value to the sponsored recommendation relative to organic recommendations. Because the retailer in our setting did not run a sponsored recommendation market, we do not have the data to precisely measure the response to sponsored recommendations. However, our approach can be applied by retailers such as Amazon and AliBaba that implement sponsored recommendation market and already have the data to conduct the appropriate analysis. Our simulation analysis is meant to show how a retailer can leverage our approach to improve efficiency of its sponsored recommendation market rather than provide a precise measure for our partner retailer. Further, because consumer response to sponsored recommendations may be attenuated relative to organic recommendations, another interpretation is

[15] We computed the revenue gains from sponsored display of product 343073 on all target focal products' pages but show only the top 10 such target products due to space constraints.

that our estimates from policy simulations provide the upper bound of potential product sales due to their sponsored display in other products' pages.

7.0 Conclusions

We conduct a large-scale randomized experiment on a fashion apparel retailer's website to estimate the causal value of recommending substitute products on the product pages. We find that while a typical product's page views increase due to its exposure on other related products' pages, its purchase probability conditional on view decreases due to recommendation of close substitutes on its own page. We further find that both location of a product in the network as well as the strength of relationship with its related products significantly affect the economic value of its recommendation links. Overall, we find that product recommendations, on average, result in a 11 percent lift in product sales. We further show that the gains do not accrue equally among products; some products with unfavorable locations in the recommendation network will lose from the widespread use of recommendation engines whereas products that are more strategically located in the network will gain substantially.

As algorithms inform more and more of consumer decisions online, the need to understand and exploit knowledge of these algorithms will become very important to managers. This understanding has already led to the emergence of a number of new marketing sub-disciplines such as search engine optimization (SEO), Amazon listing optimization and app store optimization. These specialties seek to exploit knowledge of search engine ranking algorithms, Amazon's product ranking policies and app ranking policies used by mobile app stores respectively. Similarly, our paper highlights the value to retailers and producers from investing in recommendation network optimization. Retailers like Amazon recognize this today and already allow sponsored ad networks in which producers can bid to place their products alongside other products.

The broader imperative for forward-thinking managers is that in addition to thinking of their products in terms of price, product performance, advertising and other attributes, they need to also consider the position of their products in a digital network created by recommendation engines. A product's location in such a digital neighborhood will matter on the Internet just as retail placement and location might matter in an offline world. Whether it means getting a hyperlink from a popular product on Amazon or being recommended next to Netflix's hot new documentary, there are significant gains to understanding and exploiting product recommendation networks.

We estimated the economic value of recommendation links for the retailer that we studied and hence although the signs of estimates are generalizable, their magnitudes cannot be generalized for other product categories offered by other retailers. The economic value of recommendations can be higher if they are personalized for individual customers. As different customers view the interrelationship between two products differently, and accordingly, attach different (personalized) weights for the overall co-purchase and co-view scores for product-pairs for customers based on their individual browsing behavior. This may result in different affinity scores for a link between two products in the product recommendation network for different customers – or simply personalized recommendation network for customers. We believe, such personalized recommendation links may have higher economic value and therefore exploring such personalized recommendations may be a promising area of future research. Moreover, in our field settings, the outdegree was fixed and thus we could not estimate its effect on the product demand. More studies are needed in other settings to generalize our results.

References

Adomavicius G., Z. Huang, and Alexander Tuzhilin (2008) Personalization and recommender systems. *State-of-the-art decision making tools in the information-intensive age*:55-100.

Burt R. S. (1992) Structural Holes: The Social Structure of Competition. Cambridge, Mass. Harvard University Press.

Carmi E., G. Oestreicher-Singer, U. Stettner, and A. Sundararajan (2017) Is Oprah Contagious: The depths of diffusion of demand shocks in a product network. *MIS Quarterly* 41(1):207-221.

De P., Yu (Jeffrey) Hu, and M. S. Rahman (2010) Technology Usage and Online Sales: An Empirical Study. *Management Science* 56(11):1930-1945.

Dore R. S. (1983) Goodwill and the Spirit of Market Capitalism. *British J of Sociol.* 34:459–482.

Fleder D. and K. Hosanagar, (2009) Blockbuster Culture's Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity. *Management Science* 55(5):697-712

Granovetter M. (1973) The Strength of Weak Ties. *Amer. J. of Sociol*. 78(6):1360-1380.

Kumar, A. and T. Yinliang (2015) Demand Effects of Joint Product Advertising in Online Product Videos *Management Science* 61(8):1921-1937.

Krackhardt D., N. Nohria, and B. Eccles (2003) The strength of strong ties. *Networks in the knowledge economy* 82.

Lee D. and K. Hosanagar (2016) How Do Recommender Systems Affect Sales Diversity? A Cross-Category Investigation via Randomized Field Experiment. Working Paper, 2018.

Lin Z., K. Y. Goh, and C. S. Heng (2017) The Demand Effects of Product Recommendation Networks: An Empirical Analysis of Network Diversity and Stability. *MIS Quarterly* 41(2):397-426.

Oestreicher-Singer G. and A. Sundararajan (2012) The Visible Hand? Demand Effects of Recommendation Networks in Electronic Markets. *Management Science* 48(11):1963-1981.

Oestreicher-Singer G., B. Libai, L. Sivan, E. Carmi, and O. Yassin (2013) The Network Value of Products. *Journal of Marketing* 77(3):1-14.

Podolny J. M. (2001) Networks as the pipes and prisms of the market. *Amer. J. of Sociol*. 107(1): 33-60.

Senecal S., and J. Nantel (2004) The influence of online product recommendations on consumers' online choices. *Journal of Retailing* 80(2):159-169.

Sharma, A., J. M. Hofman, and D. J. Watts (2015) Estimating the causal impact of recommendation systems from observational data. In Proceedings of the Sixteenth ACM Conference on Economics and Computation 453-470.

Uzzi B. (1996) The Sources and Consequences of Embeddedness for the Economic Performance of Organizations: The Network Effect. *Amer. Sociol. Rev.* 61:674–98.

APPENDIX A

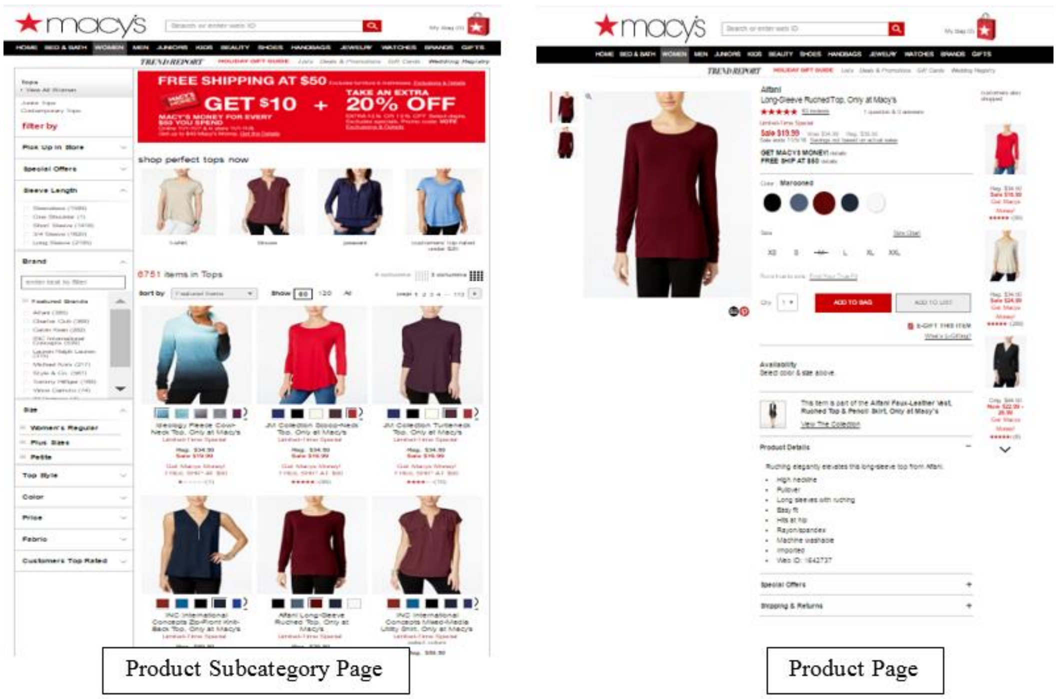

Figure A1: A recommended product appearing with the focal product on both, product subcategory main page and focal product's page

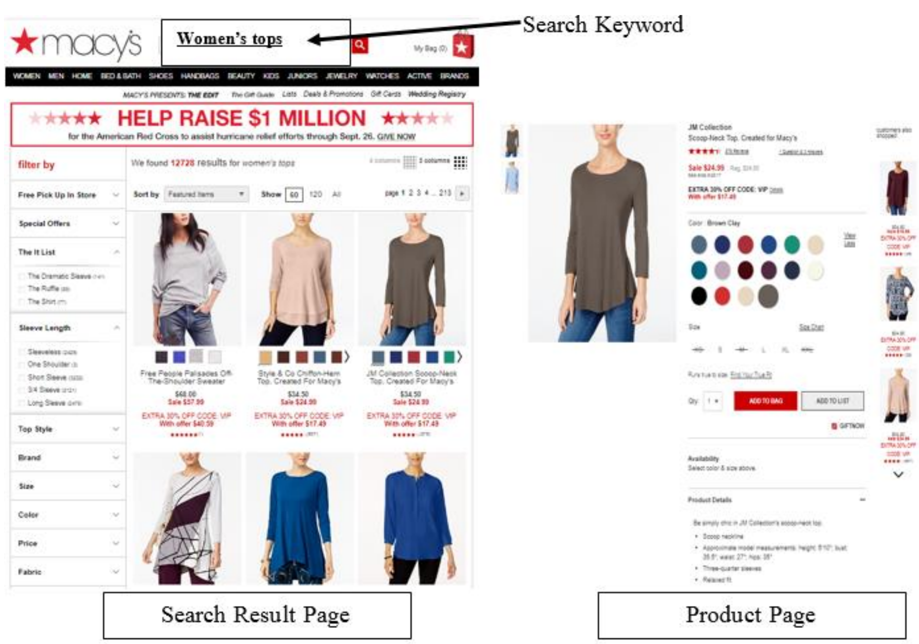

Figure A2: A recommended product appearing with the focal product on both, search result page and focal product's page.

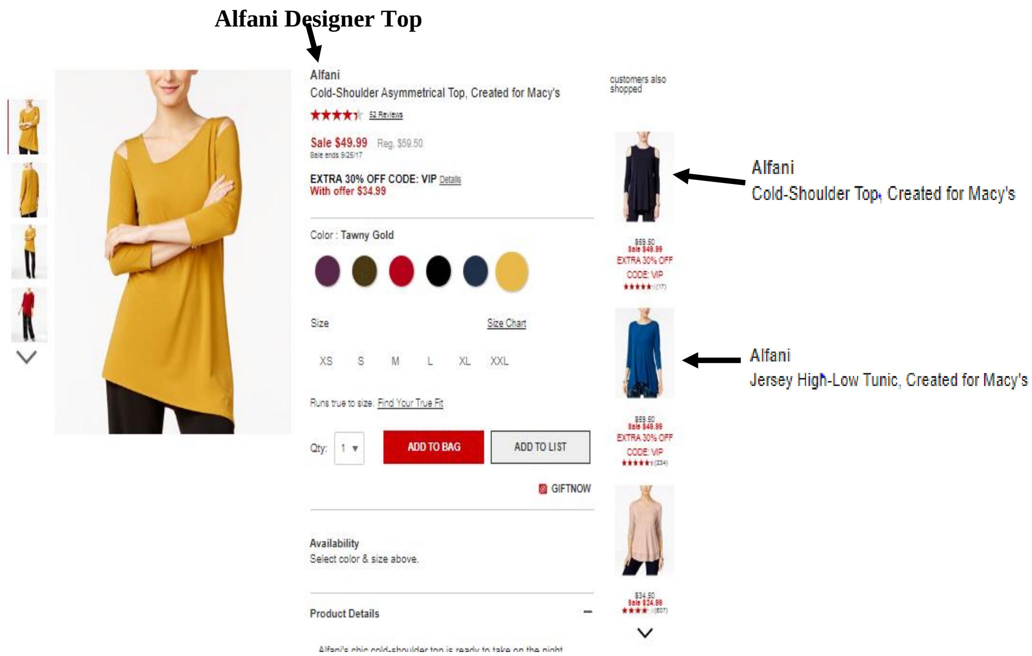

Figure A3: Recommended and focal products have similar style from the same designer

APPENDIX B

We have used Google's PageRank, the most popular eigenvector centrality measure in E-commerce applications, in our main analysis. We further conduct our analysis on all types of network centrality measures: Indegree (degree centrality); closeness (closeness centrality); weighted betweenness (betweenness centrality); and Google's PageRank (eigenvector centrality). Table B1 reports the summary statistics for mean values and range of values for these centrality measures across the 37619 products.

37619 products	Percentile values						
	0	25	50	75	90	95	100
Mean. Indegree	0	0	1.1	3.7	8.1	13	299.5
Range of indegree	0	0	4	8	16	24	452
Mean PageRank	4.6E-06	4.9E-06	7.6E-06	2.7E-05	7.3E-05	1.2E-04	1.4E-03
Range of PageRank	0	5.1E-07	1.5E-05	7.3E-05	1.7E-04	2.6E-04	2.6E-03
Mean closeness	0	0	3.0E-09	1.5E-08	3.4E-08	5.2E-08	3.1E-07
Range of closeness	0	0	1.4E-08	3.9E-08	7.6E-08	1.0E-07	4.3E-07
Mean Wgt. betweenness	0	0	2.5E-08	2.2E-07	1.0E-06	2.4E-06	6.7E-05
Range of Wgt. betweenness	0	0	1.9E-07	1.6E-06	6.9E-06	1.6E-05	2.1E-04

Table B1: Summary statistics for centrality measures

Next, we separately estimate specification (4) with each of these centrality measures to examine whether the coefficient of our interest vary with our choice of centrality measure. Table B2 reports the resulting coefficient estimates. We find that all coefficient estimates remain qualitatively similar with inclusion of different centrality measures. This shows that our results are robust to the choice of centrality measure.

Coefficient estimates (Std. Errors)	Degree Centrality (Indegree)	Closeness Centrality (closeness)	Betweenness Centrality (Wgt. Betweenness)	Eigenvector Centrality (Google's PageRank)
Dependent variable – Log of daily number of product page views				
Treat	-0.007*** (0.0005)	-0.006*** (0.0005)	-0.005*** (0.0004)	-0.008*** (0.0005)
*Treat * (Centrality Measure)*	0.0015*** (0.0001)	7.5E+05*** (2.7E+04)	2092.2*** (202.8)	346.5*** (10.31)
*Treat * Log avg. incoming affinity score*	0.013*** (0.0003)	0.01*** (0.0003)	0.014*** (0.0003)	0.011*** (0.0003)
Product-day Fixed effects	Yes			
N(No. of product-days)	3958166 (1979083)			

***, **, and * denote statistically significant at α=0.01, 0.05, and 0.10 levels (two-sided test), respectively
Standard errors cluster corrected at product-day level are in parentheses

Table B2: Regression estimates with different centrality measures

We further compute correlations between different centrality measures in Table B3, which indicates a very high correlation of PageRank with closeness and Indegree but a relatively lower value for correlation with Wgt. betweenness.

Correlations	Indegree	closeness	Wgt. Beetweenness	PageRank
Indegree	1			
closeness	0.570	1		
Wgt. Beetweenness	0.233	0.397	1	
PageRank	0.752	0.702	0.196	1

Table B3: Correlation between centrality measures

	Log daily number of product page views
Treat	-0.008*** (0.0005)
*Treat * PageRank*	347.8*** (10.5)
*Treat * Wgt. Betweenness*	1148.1*** (203.03)
*Treat * Log avg. incoming affinity score*	0.011*** (0.0003)
Product-day Fixed effects	Yes
N(No. of products)	3958166 (1979083)

***, **, and * denote statistically significant at α=0.01, 0.05, and 0.10 levels (two-sided test), respectively
Standard errors cluster corrected at product-day level are in parentheses

Table B4: Regression estimates with PageRank and Betweenness centrality

Accordingly, we estimate specification (4) with both PageRank and Wgt betweenness and report the resulting coefficients in Table B4. We find that the coefficient estimates of the treatment variable and its interaction with the PageRank remain qualitatively similar with inclusion of betweenness centrality. Moreover, the coefficient of interaction of betweenness centrality with treatment variable is also as expected positive and significant.

APPENDIX C

We examine the sessions in which products with reciprocal focal-recommended product relationship are viewed. Let's assume product A appears as recommended product on product B's page and B appear as recommended product on A's page. For such reciprocal product pair A-B, if a visitor first views product A's page, then views product B's page, and then again views product A's page. Such product page views will be included in our specification (3): once as product A's page view with recommended B's page view and once as product B's page view with recommended A's page view. If product B is purchased in this session, it will be counted as recommended product sales in the first inclusion and as focal product sales in the second inclusion in our specification (3). Specification (3) will thus identify the lift in sales of B at the expense of sales of A in the first counting and the opposite in the second counting, and thus estimate the net effect of these two effects.

	Log Product sales | product view		Log Recommended product sales | product view	
	Coeff. Est.	St Err.	Coeff. Est.	St Err.
All sessions				
Treat	-0.019***	0.002	0.09***	0.001
N (No of Product-days)	2326402 (701730)		2326402 (701730)	
Sessions with reciprocal FP-RP relationship products' page views				
Treat	-0.012***	0.003	0.05***	0.001
N (No of Product-days)	1032575 (378398)		1032575 (378398)	
Rest of the sessions				
Treat	-0.024***	0.002	0.12***	0.001
N (No of Product-days)	1293827 (323332)		1293827 (323332)	
Product-days fixed effect	Yes		Yes	

∗∗∗, ∗∗, and ∗ denote statistically significant at α=0.01, 0.05, and 0.10 levels (two-sided test), respectively
Standard errors cluster corrected at product-day level are in parentheses

Table C1: Estimates with reciprocal FP-RP relationship

To check whether our results are robust to such simultaneity issue, we estimate the effect of recommendations for sessions in which focal products with reciprocal focal-recommended products relationship were viewed and rest of the sessions separately. Table C1 reports the resulting estimates as well as the estimates with all sessions for easy comparison. We find that our results are robust to exclusion of sessions with page views of products with reciprocal relationship.

APPENDIX D

We examine whether the effects of recommendations are different across different product categories. The retailer in the present field setup sells over 35,000 different products under the following broad product categories: Home goods (such as luggage, home décor, outdoor and recreational goods, and kitchen and dining); Bed and bath goods (such as bedding, pillows, quilts, bath rugs, curtain, towels and accessories); Women's apparel (such as tops, dresses, pants, jeans, skirts, skorts, shorts, and swimwear); Women's lingerie and sleepwear; Men's apparel and accessories (such as shirts, t-shirts, pants, jeans, shorts, ties, vests and briefs); Apparel and accessories for teenage boys and girls; Baby apparel, toys, and accessories; Shoes and sandals; and Accessories (such as fragrance, cosmetics, handbags, and jewelry). Products in some categories, such as apparel (for men, women, and teenagers), accessories, and shoes & sandals fall under the fashion domain (more hedonic in nature). Products in other categories, such as home goods, bed & bath products, and baby products (baby apparel, toys and accessories) do not fall under fashion domain. Accordingly, we estimate specification (2) and (3) separately for products that fall under and do not fall under fashion domain. We report the coefficient estimates for the full sample and these two subsamples of products in Table D1, which suggests that our results hold for both categories of products.

	Log No. of daily product page views		Log Product sales | product view		Log Recommended product sales | product view	
	Coeff. Est.	St Err.	Coeff. Est.	St Err.	Coeff. Est.	St Err.
Full Sample of products						
Treat	0.075^{***}	0.001	-0.019^{***}	0.002	0.09^{***}	0.001
N (No of Product-days)	3958166 (1979083)		2326402 (701730)		2326402 (701730)	
Products in fashion domain						
Treat	0.098^{***}	0.001	-0.027^{***}	0.002	0.09^{***}	0.001
N (No of Product-days)	2719208 (1359604)		1964605 (524917)		1964605 (524917)	
Products not in fashion domain						
Treat	0.05^{***}	0.004	-0.011^{**}	0.005	0.08^{***}	0.003
N (No of Product-days)	1238958 (619479)		361797 (176813)		361797 (176813)	
Product-day fixed effect	Yes		Yes		Yes	

∗∗∗, ∗∗, and ∗ denote statistically significant at α=0.01, 0.05, and 0.10 levels (two-sided test), respectively
Standard errors cluster corrected at product-day level are in parentheses

Table D1: Estimation results for different product categories