JACOB GOLDENBERG, GAL OESTREICHER-SINGER, and SHACHAR REICHMAN*

Online content and products are presented as product networks, in which nodes are product pages linked by hyperlinks. These links are typically algorithmically induced recommendations based on aggregated data. Recently, websites have begun to offer social networks and user-generated links alongside the product network, creating a dual-network structure. The authors investigate the role of this dual-network structure in facilitating content exploration. They analyze YouTube's dual network and show that user pages have unique structural properties and act as content brokers. Next, the authors show that random rewiring of the product network cannot replicate this brokering effect. They present seven Internet studies in which participants browsing a YouTube-based website are exposed to different conditions of recommendations. The first set of studies shows that exposure to the dual network results in a more efficient (time to desirable outcome) and more effective (average product rating, overall satisfaction) exploration process. The next set of studies extends the previous ones to include dynamic structures, in which the network changes as a function of time or in response to participants' satisfaction. Furthermore, the results are replicated using data from another content site.

Keywords: e-commerce, social networks, product network, user-generated content

The Quest for Content: How User-Generated Links Can Facilitate Online Exploration

With the seemingly limitless variety and quantity of content available online, consumers devote substantial effort to searching for content and products that match their tastes. The growth of online content has recently been bolstered by the increasing popularity and prevalence of user-generated content (UGC) and Web 2.0 technologies. Considering that online content is rarely accompanied by traditional marketing campaigns, two fundamental questions in this context arise: How do consumers find "good" content? and How can the site owner (e.g., Amazon.com, YouTube) increase the efficiency (i.e., the time it takes the consumer to find satisfactory content) and the effectiveness (i.e., the consumer's satisfaction from the search results) of the consumer's search process?

Consider the following example. On September 13, 2009, the famous artist Kanye West surprised the world as he jumped on the stage of the MTV awards, grabbed the microphone from the artist Taylor Swift (who was receiving the award for "Best Female Video of the Year"), and stated that another performer deserved the award more. On the same day, Yarin Lidor posted to YouTube a video mocking the event, titled "Remix: Kanye West Disrespect and Interrupt Taylor Swift." Despite the creator's anonymity, the video was a great success and was viewed more than three million times over the following two weeks, making it one of the

most viewed videos of those weeks. Figure 1, Panel A, provides the summary statistics of the distribution of sources from which those three million viewers (the "consumers") arrived at the video's page. Because Kanye West's outburst received much media attention, a large portion (approximately 39%) of the consumers arrived at the page after a keyword search on this topic. Approximately 12% arrived at this page after following a "related video" link on another video's page, and 22% arrived after following a link placed on a user's personal page on YouTube.

Six days later, Mr. Lidor posted another video on YouTube, named "How Could You Be So Jackass." That video received milder success and was viewed by approximately 30,000 viewers. Figure 1, Panel B, provides the summary statistics of the distribution of sources from which those consumers arrived at the video's page. Clearly being less popular, this video was not searched for as often (only 1% of consumers arrived at the page after using a keyword search). However, approximately 40% arrived at this page after following a "related video" link on another video's page, and 19% arrived after following a link placed by a user on the YouTube social network.

The traffic patterns to both of Mr. Lidor's videos suggest that the "related videos" (i.e., YouTube recommendations)

Figure 1
DISTRIBUTION OF THE SOURCES OF TRAFFIC

A: Yarin Lidor's First Video

B: Yarin Lidor's Second Video
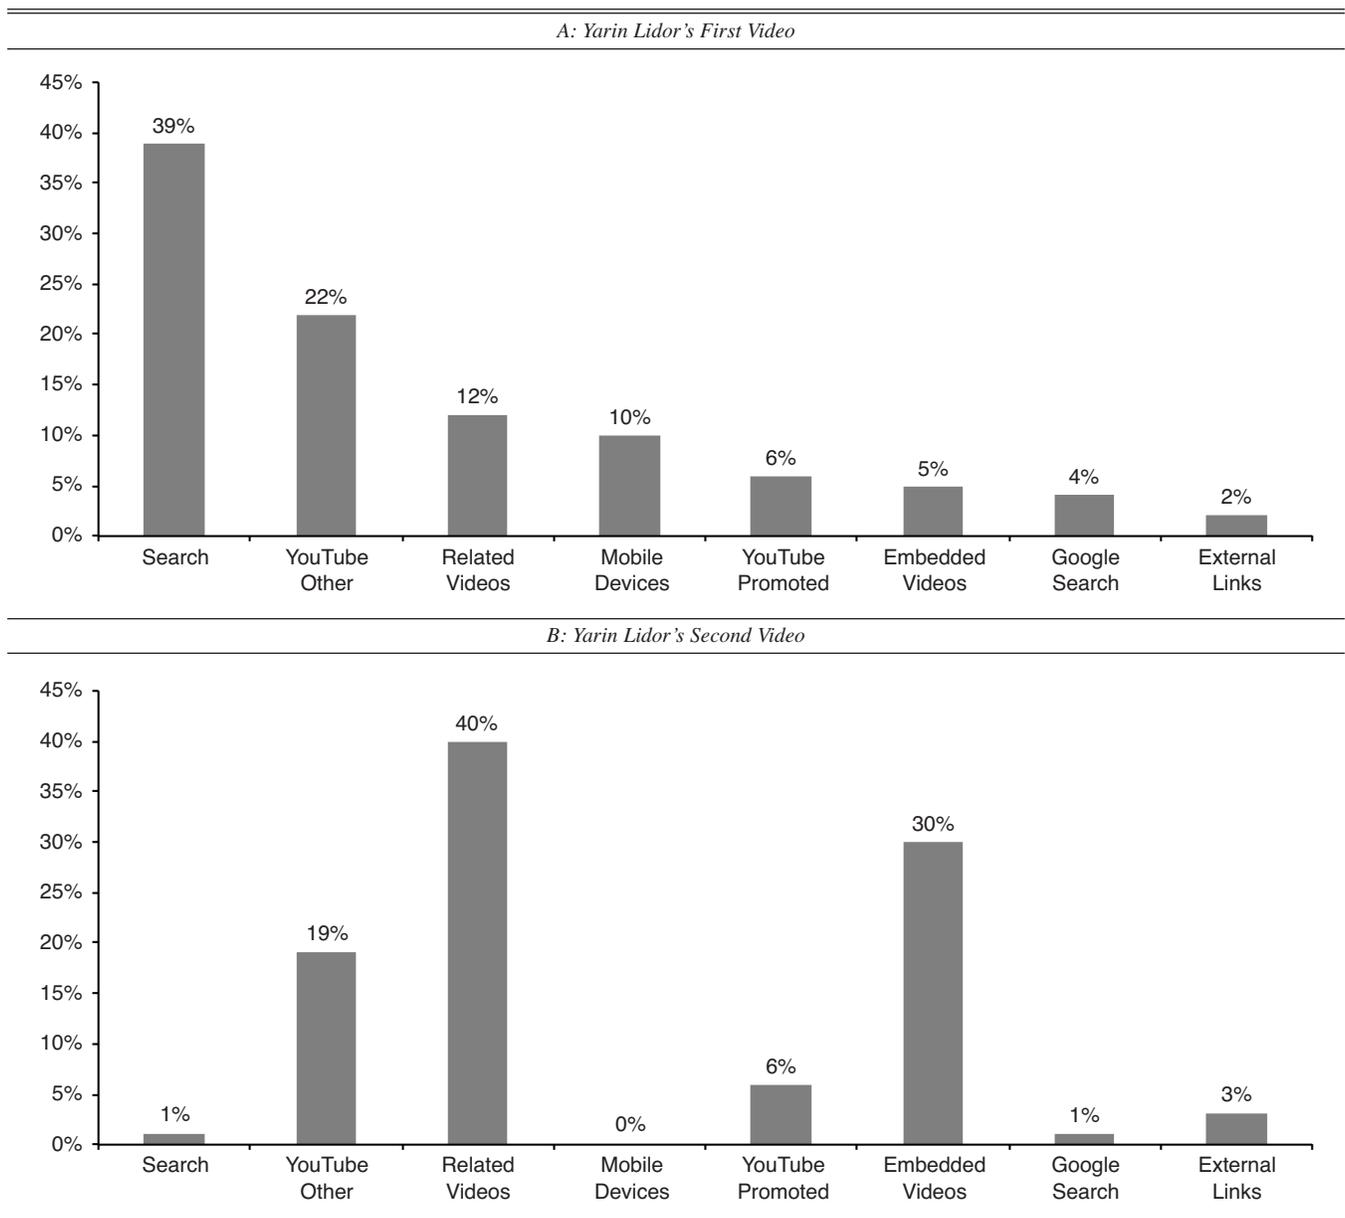
Notes: Each column shows the percentage of users who arrived at each video from a specific source as follows: "Search" refers to users who found the video using a keyword search in YouTube, "YouTube Other" refers to users who arrived at the video from a user page, "Related Videos" refers to users who arrived at the video after clicking on a "related video" link on another video's page, "Mobile devices" refers to users who use mobile devices, "Embedded Videos" refers to users who followed a link to the video embedded on a web page other than YouTube.com, "Google Search" refers to users who found the video using a keyword search on Google, and "External Links" are users who followed a direct link from other sites.

and links placed by users on the YouTube social network both play important roles in directing consumers' attention. They also highlight another salient characteristic of consumption of UGC: the existence of content exploration with no clearly defined target. Of the viewers who viewed Mr. Lidor's first and second videos, 34% and 59%, respectively, found the content with no direct intention to view it specifically, but rather through exploration based on recommendations.

While the first video received media attention, most content on UGC sites is more similar to the second video. That is, few consumers specifically search for them, and most views are the result of exploration, also referred to as hedonic browsing (Moe 2003; Moe and Fader 2004). In this article, we study how user-generated links affect exploration efficiency and effectiveness. We focus on consumers who explore the content space with no defined target and are therefore less likely to use traditional keyword search.

A structure that content and retail websites have widely adopted is the product network. In this structure, each product is connected by links to other products' web pages. The links are usually algorithmically induced recommendations that map complementarity relationships between products (e.g., copurchased products, such as Amazon.com's recommendations appearing under the title "Customers who bought this item also bought…"). The product network structure allows consumers to browse a wide variety of content and has been shown to influence their consumption patterns (Oestreicher-Singer and Sundararajan 2012a).

Another, better-known, type of network prevalent online is the social network, which typically describes relationships between individuals (e.g., friends, colleagues, trading partners). These networks can be described as collections of personal web pages that are linked to one another, in which a link represents a connection between two individuals. Social networks are prevalent mainly in designated websites, such as Facebook or LinkedIn. Recently, a growing number of e-commerce sites that use product-network-based recommendation systems have begun to add social networks. On such a site, the product network is linked to the social network through user-generated links: links between users' pages and products' pages. These links typically appear in the context of online opinions or conversations (i.e., comments, reviews, or posts). Each link is created by a single individual without any intervention of a recommendation algorithm or aggregate data and thus represents that person's own opinion about the product. This link can be thought of as a personal recommendation from that individual regarding the product. Note that in contrast to two-mode networks, edges form between the same types of nodes (i.e., social and product network) as well as between different types of nodes (Buldyrev et al. 2010; Newman 2003). The integration of the social network with the product network creates a new structure, which we refer to as the "dual-network structure" (illustrated in Figure 2). Dual networks are becoming prevalent on UGC sites, such as YouTube, Flickr, and Digg. Some e-commerce sites have slowly begun to incorporate social features such as communities and forums in which consumers can actively participate. For example, on Amazon.com, a customer book review on a product page can include a link to the reviewer's personal page, and a reviewer's personal page might include links to other books or other reviewers.

Figure 2
ILLUSTRATION OF THE YOUTUBE.COM DUAL NETWORK
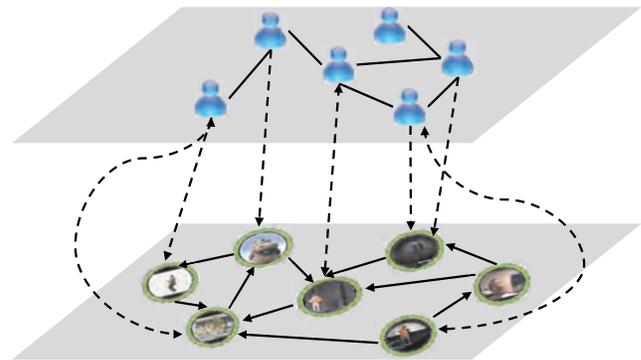
Notes: The dotted arrows represent user-generated links between social network nodes (denoted as icons) and product nodes (denoted as circles), solid lines represent social links between users, and solid arrows represent links between products (i.e., "Related Videos" links).

We focus on the role of this dual-network structure and specifically of user-generated links in facilitating content exploration. We analyze the YouTube dual network as an example for online content sites. Using data on more than 700,000 videos and users from the YouTube dual network, we are able to show that user pages have unique structural properties (e.g., high betweenness centrality) that enable them to act as better content brokers in the dual network. Next, we show that random rewiring of the product network cannot replicate the brokering effect of the user-generated links of the real-world dual network. We introduce seven web studies in which consumers browse a YouTube-based website and are exposed to different conditions of recommendations based on the YouTube product network, the YouTube dual network, and a randomized network we created. Our first four studies show that allowing consumers to use both the product network and user-generated links leads them to find satisfactory content more quickly. We then replicate our study using a dual network from a different website—Last.fm—and obtain similar results. We subsequently extend these studies to include a dynamic strategy, offering stronger managerial implications for e-commerce policy. In these studies, the types of links consumers are exposed to change as a function of time (Studies 5 and 6) or in response to participants' actions (Study 7). We find that these dynamic structures lead to an increase in consumers' ratings of the content and to higher overall satisfaction.

LITERATURE REVIEW

Most research on online content search focuses on keyword search (Arasu et al. 2001). However, as Rangaswamy, Giles, and Seres (2009) indicate, search engines still fail to meet the needs of users who explore the content space without a clear idea of what they are looking for. Marketing scholars have noted the importance of the distinction between directed and exploratory search, and they have become increasingly interested in online shopping behavior, aiming to better understand the path that leads to conversion (for a framework for such research, see Hui, Fader, and Bradlow 2009). Moe (2003) presents a taxonomy of behav-

ior for online shoppers on a retail site, classifying shoppers according to their motivations and goals. More recently, Moe and Fader (2004) studied the conversion rate in online shopping, distinguishing visits motivated by a planned purchase and visits associated with hedonic browsing (see also Bucklin et al. 2002). We focus on exploratory search, or hedonic browsing, in the context of search for content.

Many content websites aim to guide consumers' exploration processes (e.g., to improve users' experience of the site) or to provide exposure to certain types of content. In general, websites guide consumers by exposing them to recommendations generated by designated algorithms. Given the recent advances in recommendation algorithms (e.g., Adomavicius and Tuzhilin 2005; Herlocker et al. 2004), why would user-generated links (and dual-network structures) contribute to the process of exploration over a product network? The answer may be found through analogy to the social network. Recent marketing literature has connected network structure properties to information dissemination in social networks (e.g., Goldenberg et al. 2009; Goldenberg, Libai, and Muller 2001; Katona, Zubcsek, and Sarvary 2011; Shaikh, Rangaswamy, and Balakrishnan 2006; Trusov, Bucklin, and Pauwels 2009; Valente 1996; Van den Bulte and Wuyts 2007). Specifically, the informative advantage of weak ties and structural holes in social networks has been demonstrated repeatedly (beginning with Burt's [1992] and Granovetter's [1973] seminal works). In principle, people who have strong ties usually possess similar information and have little to contribute to one another (e.g., in a job search scenario, as in Granovetter [1973]). In contrast, people with weak ties belong to different social circles, and therefore information they transmit to one another is more likely to be new.

In the context of the product network, recommendation systems are likely to link content items that are broadly similar to one another, with relatively low variance. This occurs for two reasons. First, such systems are heavily built on data that come from directed (keyword) search. When a person searches for specific content (e.g., "*Seinfeld* episode"), his or her subsequent searches are likely to be for similar content (e.g., more episodes of *Seinfeld*, interviews with the actors). Second, the recommendation system aggregates the choices of the entire population with the goal of achieving a high overall likelihood of satisfaction. The intersection of such preferences is likely to exhibit relatively low variety in content. The result of these two mechanisms is that, in most cases, the algorithm recommends products that are similar to one another. For example, we find that on YouTube, 56% of the product network links connect videos within the same category.

In social network research, frequency of communication (e.g., number of telephone calls, number of joint posts) is often used as a proxy for tie strength. Applying this logic to the product network setting, we can consider recommendation links as strong ties between products, given that these links are based on the preferences of the majority (e.g., being frequently co-viewed or co-tagged). We explore and confirm this premise in the empirical part of this article. As we discussed previously, during the exploration process, and especially as time elapses, strong ties may cease to be interesting or informative because they point to similar content (e.g., videos, books). To branch out from the close circle of similar content, weaker ties are required.

These weak ties are provided by the dual-network structure. Just as, according to social networking theory, a person can bridge across several social circles, we can view each individual user of a website as being situated in the intersection of different types of product circles (e.g., different genres, different categories). The typical user does not post links to identical content types but rather presents a variety of content to reflect his or her different preferences. For example, a user may have favorite funny commercials as well as favorite educational videos, which have little in common. Thus, users will often bridge across products of different circles according to their preferences (e.g., on YouTube, less than 20% of a given user's generated links connect products of the same category). The resulting user-generated links function as the weak ties of the product network. Because content exploration is an ongoing process, which often includes seeking new information, it is possible that users will be able to benefit from the ability to observe weak ties (user-generated links) together with the strong ties (algorithm-induced recommendations).

To the best of our knowledge, our work is the first to study the dual network often present in electronic commerce sites. Work on product networks includes a study of the network of networked content sites by Katona and Sarvary (2008), who investigate how to strategically link between websites in a market of advertising links. However, product networks were not explicitly mentioned in that study. Oestreicher-Singer and Sundararajan (2012b) study the network of books on Amazon.com and quantify the incremental correlation in book sales attributable to product networks' visibility. We propose that the integration of social and product networks will facilitate exploration of content and that the exploration process is more efficient in a dual-network structure that incorporates user-generated links than in a regular product network based on algorithmically induced recommendations.

OVERVIEW OF DATA

Using data from YouTube, one of the largest existing dual networks, we conduct an in-depth analysis of the dual-network structure. YouTube's core business is centered on videos, which are the website's "products." Each video has an associated web page that is connected by links to other videos' web pages, thus creating a product network. In addition to the product network, YouTube offers a social network in which each user has an associated web page (Figure 3 presents sample product and user pages), and these webpages can be linked to other user pages (creating a social network) and to video pages (connecting the product and social networks). This creates a dual-network structure (see Figure 2).

Using a crawler, we collected data on the YouTube product network and social network. We collected data for approximately 700,000 videos and for 50,000 users connected by approximately 10 million links. The links include algorithm-generated links between videos based on co-consumption (labeled "Related Videos"), social links between user pages ("Friends"), and user-generated links between products and users ("Owner" and "Favorites").

Figure 3
ILLUSTRATIONS OF A PRODUCT (VIDEO) PAGE AND A USER PAGE ON YOUTUBE

A: Product (Video) Page	*B: User Page*

Notes: "Friends" are links in the social network, "Related Videos" are links within the product network, "Favorites" and "Own Videos" are user-generated links from a social node to the product network, and the "Video's Owner" link is a user-generated link from a product node to the social network.

Data were collected using snowball network sampling, which is a common technique used in large network sampling (Ahn et al. 2007; Carrington, Scott, and Wasserman 2005; Wasserman and Faust 1994). Specifically, we used a breadth-first algorithm beginning with the 25 most viewed videos on YouTube.com, following each video's links to its owner and to its related videos. We then followed user friend links as well as related-video links and continued up to four hops from the source (fourth neighbor). At the fourth level of the data collection, we collected only outgoing links of nodes (videos) to other nodes already in our data set. Incomplete information is a common problem when sampling a network of this size, especially regarding the outer edges of the network. By using this method, which collects data on the links of these outer nodes to inner nodes in the sampled network, we reduce the level of incompleteness.

THE TOPOLOGY OF DIFFERENT NETWORK STRUCTURES

The first step in our empirical investigation required in-depth analysis of the dual network and a comparison of this structure with the two alternative network structures: the product network and the synthetic dual network. To create an example of a product network, we extracted only the product nodes and the links between them from the previously mentioned data set. To make sure all the networks that we considered were of comparable sizes and because there were 50,000 user nodes in the dual network, in the studies described subsequently, we used a reduced version of the dual network, which we generated by eliminating 50,000 randomly chosen product nodes from the outer edge of the original dual network.

The dual-network structure may facilitate a more efficient exploration process. Therefore, it is important to understand whether this is due to the unique information carried over user-generated links or simply a result of offering a more diverse set of options, in which case the same benefits could be achieved easily by implementing a more advanced recommendation algorithm. The literature on network structure (Newman 2003) suggests that adding random links reduces the average distance between products in a network. Therefore, such rewiring could potentially assist the exploration process as well. To be able to compare such a network with the dual network, we constructed a synthetic dual network.

This network is based on the YouTube product network mentioned previously (the reduced network with 50,000 user nodes eliminated), with 50,000 artificially created dummy-user pages, each randomly connected to different products, according to the degree distribution of the real user nodes.

Using the three networks (the YouTube dual network, the YouTube product network, and the synthetic dual network), we computed several indexes that are commonly used in the literature to characterize network structures and effectiveness (Newman 2003; Wasserman and Faust 1994), including the following (for detailed definitions of the indexes, see Appendix A):

1. Degree (number of links) of each node, including the in-degree (the number of incoming links) and out-degree (the number of outgoing links);
2. Closeness centrality of each node: a measure of the average minimal distance (number of hops) between this node and any other node in the network;
3. Betweenness centrality (Freeman 1977) of each node: a measure of the number of shortest paths (between any two nodes) in which this node is included;
4. PageRank of each node (Brin and Page 1998): an iterative measure of centrality, which is based on the number of links pointing to a node and the centrality of those links; and
5. Assortative mixing level of each node: a measure of the level of similarity (homophily) between a given node and its neighbors (Newman 2003).

Note that we treated all types of nodes as one integrated network when computing the indexes. Table 1 shows our findings for the YouTube dual network and for the synthetic dual network, separated according to the types of nodes.

Observing the YouTube dual network (the two left columns of Table 1), it is perhaps surprising that even when both types of nodes are included in one combined network, they have notably different structural properties. Most important, compared with video (product) nodes, user nodes in the YouTube dual network have a significantly higher betweenness (more than three times higher). This observation suggests that user pages (and their associated user-generated links) play an important role in increasing network connectivity. These results are even more salient when contrasted with the results involving closeness centrality and the nodes' average degree.

Video nodes have a higher closeness centrality (low distance), meaning that video nodes are, on average, closer than user nodes to any other video in the network. This is consistent with the observation that video pages have a

higher in-degree and out-degree than user pages, probably due to lower levels of activity in the social network compared with the product network. If the product nodes and user nodes served the same type of function, such differences would lead to a higher betweenness centrality among product nodes than among user nodes. However, the opposite is the case: The betweenness centrality is more than three times higher for user nodes than for product nodes. Taken together, these findings imply that user nodes play an important role in bridging different parts of the network and possibly facilitating faster exploration of the content space. This suggests that the user-generated links (links connected to the user nodes) are different from algorithm-induced nodes of the product network, and the combination of the two networks might lead to a superior structure for searches.

In accordance with the literature on random rewiring, the synthetic network shows a significantly lower average distance between products (the two rightmost columns of Table 1). However, the betweenness of the user pages in the dual network is approximately two times that of the added dummy-user pages (i.e., random nodes) in the synthetic network. These results suggest that the user-generated links encompass implicit characteristics that cannot be easily imitated algorithmically. Users seem to be able to create structural variety on their pages and group various videos together, reflecting the idea that different videos may be related according to a particular user's perception.

The results regarding assortative mixing are of special interest in this context. In the study of complex networks, the term "assortative mixing" refers to a bias in favor of connections between network nodes with similar characteristics (also referred to as "homophily" in the context of social networks). We adjust the assortative mixing measure to our context and define it as follows: For a product (i.e., video), the level of assortative mixing is the percentage of linked videos ("Related Videos") that are of the same category; for a user, the level of assortative mixing is the percentage of linked videos ("Favorites") that are of the same category. As Table 1 shows, the assortative mixing of a video is, on average, 56%, meaning a video is most likely to be linked to other videos that share its category. Users' assortative mixing is, on average, 19%, meaning that only approximately one-fifth of a given user's favorite videos are of the same category; that is, user-generated links seem to bridge between categories more than 80% of the time. This further supports the finding of the important role of users' pages as brokers in the product network. Next, we examine the direct influence of user-generated links on exploration-based search efficiency. Can such links "land" video seekers on new content with more meaningful variety?

DUAL-NETWORK EFFECT ON EXPLORATION: WEB-BASED STUDIES

In light of our empirical observations, the question that arises is whether the dual network can affect consumers' consumption. Do the structural properties of this network indeed affect the efficiency (time needed to find sufficiently interesting content) and effectiveness (overall satisfaction from the content) of the exploration process?

Table 1
INDEXES FOR THE DUAL NETWORK AND RANDOMIZED LINKS, DIVIDED ACCORDING TO THE TYPES OF NODES

| Network Indexes | The Dual Network | | Synthetic Network | |
	Product Pages	User Pages	Product Pages	Dummy User Nodes
Closeness	6.31	6.64	6.06	5.95
Betweenness	3.59×10^6	13.6×10^6	3.16×10^6	7.73×10^6
InDegree	14.73	4.71	14.06	4.21
OutDegree	14.70	5.20	13.96	4.32
PageRank	.99	.70	1.00	.30
Assortative mixing	.56	.19	.56	.11

Overview of the Studies

To examine the role of the dual network, we designed a set of seven web studies on a YouTube-like website specially designed and built for this study. The website provides video pages that can be viewed using a built-in video player. Each page contains ten links to other videos, enabling navigation through the site. The site offers more than 500,000 videos and recommendations collected from YouTube. For a person who browses this site, the appearance and feeling are identical to that experienced browsing YouTube.

In Study 1, we show that when consumers are allowed to use both recommendations from the product network and user-generated links from the social network, they find satisfactory content faster than when they are restricted to use of the product network or the user-generated links alone. We replicate this finding in three different setups, and we also replicate this study (Study 1a) using a dual network from a different website—Last.fm—and obtain similar results.

In Studies 2–4 we rule out a variety of alternative explanations for the superior performance of the dual network. In Study 2, we rule out the explanation that the dual network is superior merely because it adds variety through its recommendations; in Study 3, we rule out the possibility that the dual network merely points to content of higher quality; and in Study 4, we rule out an alternative explanation that labeling, and specifically the personal nature of the generated links, drives the results of Study 1. Then, in Studies 5–7, we show that use of a dynamic recommendation strategy, in which users are exposed to different types of links at different times, significantly increases the effects of the linked network.

Study 1: The Effect of the Dual Network on Exploration-Based Search

Method. To examine the effect of the network structure on exploration efficiency and effectiveness in Study 1, we first presented participants with different sources of recommendations, enabling recommendations from the product network alone, from user-generated links alone, and from the dual network. Six hundred twenty participants (48% female, mean age = 25.9 years) watched videos for approximately 15 minutes and assigned a rating from one star ("Poor") to five stars ("Awesome!") to each video they watched. (This was the same rating system YouTube was using at the time of data collection.) After watching one video, the participant was free to choose the next video from a list of recommendations and was subsequently presented with the video of his or her choice (with its corresponding recommendation links). We designed the entire navigation experience to be exactly as on YouTube. Tasks were submitted as human intelligence tasks to Amazon Mechanical Turk. Every participant was asked to browse the study web page corresponding to the assigned scenario. Participants were required to complete the entire task, watching videos for approximately 15 minutes, to receive payment. Each participant received $1.50 upon completing the task. We excluded from our analysis participants who watched fewer than 4 videos and those who watched more than 25 videos, assuming neither group could be counted as representative of exploration. Of the 620 participants, we removed 15 on the basis of these criteria (2.4%).

We randomly assigned each participant to one of three possible versions of this website:

1. Product (video) network (condition C1.1): Users of the website were offered, alongside each video they viewed, ten recommendations based on the YouTube product network only ("Related Videos" as suggested by YouTube). Those recommendations were presented under the label "Related Videos." Note that we did not alter YouTube's recommendations, and therefore they reflect the results of the site's recommendation algorithm. Figure 4, Panel A, presents an illustration of this scenario.
2. Dual network (condition C1.2): Users of the website were offered two types of recommendations, both based on the YouTube dual network. Under the label "Related Videos," they were offered five recommendations based on the YouTube product network, and under the label "People who liked this also liked," they were offered five recommendations of videos marked as "favorites" by a single user who had also marked the featured video as a favorite. We randomly selected this user from all users who marked the presented video as a favorite. Figure 4, Panel B, presents an illustration of this scenario.
3. User-generated links only (condition C1.3): Users of the website were offered ten recommendations only on the basis of user-generated links. That is, the recommendations are videos marked as favorites by a single user who had also marked the featured video as a favorite. This user was randomly selected from all users who marked the presented video as a favorite. All recommendations were presented under the label "People who liked this also liked." Figure 4, Panel C, presents an illustration of this scenario.
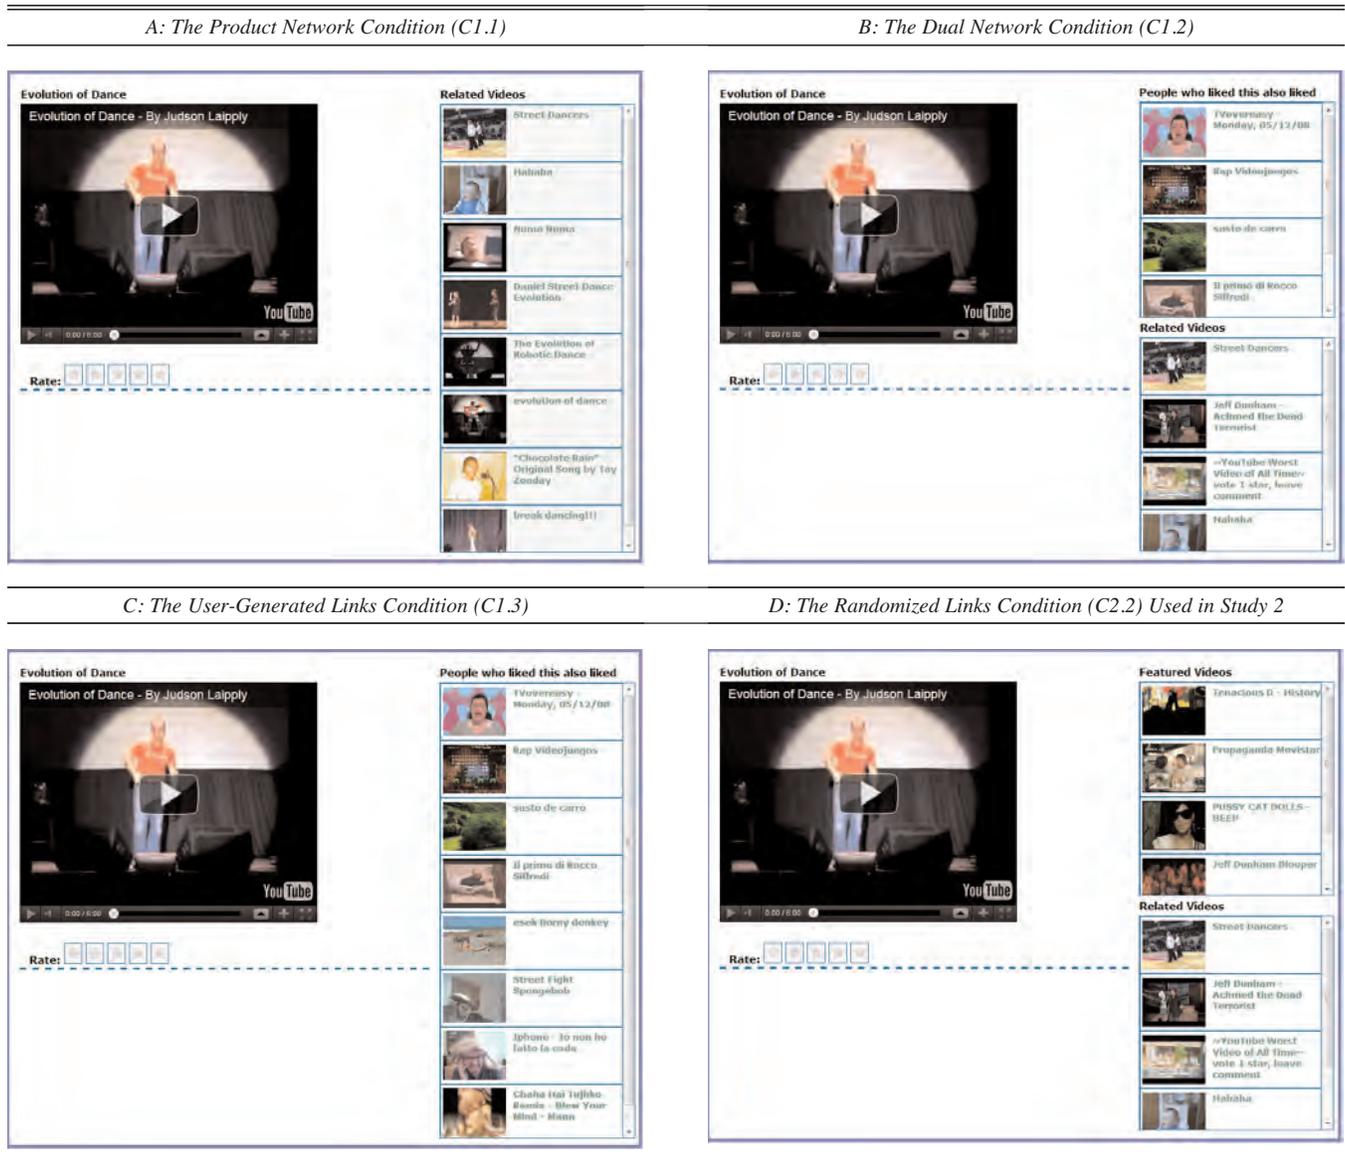
Analysis and results. We studied the effects of the different conditions (i.e., different networks) on the successful conclusion of an exploration process. We followed duration model theory and used a hazard-rate model in which the hazard rate $h(t)$ was the probability of finding "good" content. In the context of this work, the "event" occurred when a user rated a video as four or five stars (out of five) for the first time.

We first estimated a nonparametric Kaplan-Meier maximum likelihood estimation of the survival function (Kaplan and Meier 1958). Figure 5 shows the Kaplan-Meier curves for the three conditions. The curve of the dual-network condition is consistently lower than those of the product-network and the user-generated-links-only conditions. (The hazard ratios are 1.283 for the dual-network condition and .886 for the user-generated-links condition, where the product-network condition serves as a baseline hazard, $p < .05$.) That is, the time to the event (finding a liked video) seems to be shorter when using the dual network. Moreover, as the exploration time increases, the curves become even further apart, suggesting that the benefits of the dual network over the other structures are greater over time.

On the basis of these results, we defined the hazard function $h(t)$ as the instantaneous potential per second for the event (rating of four or five) to occur, given that it has not happened up to time t (Kleinbaum and Klein 2005):

$$(1) \qquad h(t) = \exp\left\{ \alpha_0 + \sum_{j=1}^{J} \alpha_j \text{Treatment}_{ij} \right.$$

$$\left. + \sum_{k=1}^{K} \beta_k \text{Demographic}_{ik} + \sum_{l=1}^{L} \gamma_l \text{VideoInfo}_{il} \right\}.$$

Figure 4
SCREENSHOTS OF THE CONDITIONS

A: The Product Network Condition (C1.1)	B: The Dual Network Condition (C1.2)

C: The User-Generated Links Condition (C1.3)	D: The Randomized Links Condition (C2.2) Used in Study 2

The Treatment variables included a dummy for each condition, excluding the product network condition, which served as the baseline. The demographic data included gender, age, and self-reported level of familiarity with the YouTube website. The video information included the rating of the first video presented to the user (rated by the general YouTube audience) and a dummy for whether it was popular (watched more than 30 million times).

We used the Cox proportional hazards model (Cox 1972) to estimate the effect of the covariates on the hazard rate, as shown in Table 2. Column A shows the effect of the different conditions, column B adds demographics of the participants, and column C adds information about the YouTube video that began the exploration process. The last estimation is based on the notion that the source of the exploration process may influence the time it takes to find "good" content.

All estimations show that use of the dual network is associated with an increase in the odds ratio under all model

Table 2
STUDY 1: PROPORTIONAL HAZARDS REGRESSION (COX REGRESSION) RESULTS

	(A) B (SE)	(B) B (SE)	(C) B (SE)
Condition			
Dual network (C1.2)	.25** (.09)	.26** (.1)	.21* (.13)
UGL only (C1.3)	−.12 (.13)	−.12 (.13)	.12 (.13)
Demographics			
Gender		−.09** (.1)	−.13* (.11)
Age		.0 (.01)	.0 (.01)
Familiar with YouTube		.01 (.03)	.02 (.03)
VideoInfo			
Start from popular video			.12 (.1)
First video rating			.19 (.27)
–2 Log-likelihood	6201.7	6197.1	6464.8

*Significance at .05 level.
**Significance at .01 level.

Figure 5
STUDY 1: KAPLAN-MEIER SURVIVAL FUNCTION RESULTS
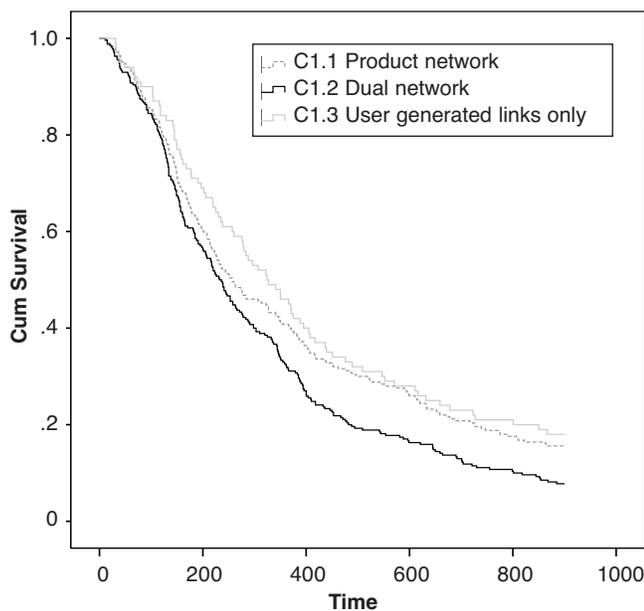
Notes: Each line represents the users of a given scenario as detailed by the legend.

specifications (the baseline model being the product-network condition). That is, participants under the dual-network condition find "good" content faster. Specifically, the hazard for the dual-network participants is approximately 30% higher than the hazard for the product-network group. Notably, exploration efficiency in the network containing user-generated links only is no better than in the product network and is worse than in the dual network. Although many of our control variables (including age, video information, familiarity with YouTube, and the rating of the first video) are not statistically significant, we find that gender has an effect on the hazard (Columns B and C). We have also replicated this analysis including an interaction term, but with no significant results.

The Effect on Overall Exploration Effectiveness

The hazard model focuses on the time until the first liked video (exploration efficiency), which can be thought of as a purchase, or conversion of the exploration process. It is important to note that many content websites' revenues are based on advertising, and it may be in their best interest to engage the user with the website for a longer period of time. Thus, we are also interested in studying the overall satisfaction of the user as well as his or her satisfaction over the course of the exploration. Therefore, we computed two additional variables for each user: AverageRating, the average video rating the focal user gives, and HighRankingScore, the proportion of highly ranked videos (ratings of four or five stars) of all videos the focal user watched.

Using a post hoc analysis of variance (ANOVA; Fisher's least significant difference test), we found that the AverageRating of participants in the dual-network condition was significantly higher than that of participants in the product-

network and user-generated-links conditions (2.95, 2.74, and 2.73 respectively; $p < .05$). As shown in Figure 6, Panels A and B, we obtained similar results for HighRanking-Score: Participants in the dual-network condition had a higher proportion of high rankings than did participants in the other two conditions ($p < .05$ using post hoc ANOVA). We observed no significant difference in either the Average-Rating or the HighRankingScore between participants in the product-network condition and those in the user-generated-links condition.

Participants also rated their overall experience with the website (on a ten-point scale). Comparing those ratings across conditions, we found that users under the dual-network condition had a significantly higher overall satisfaction rating (6.91) than users in the other two conditions (6.01 under the product-network and 5.92 under the user-generated-links condition, $p < .01$). Figure 6, Panel C, presents these results.

For robustness, we extended the analysis to include ordinary least squares regression to estimate the effect of the treatments as well as other observed covariates on each of the following dependent variables: AverageRating, HighRankingScore, and overall satisfaction. We estimated the following equations:

i. $\text{AverageRating}(i) = \alpha_0 + \Sigma_{j=1}^{J}\alpha_j\text{Treatment}_{ij}$
$+ \Sigma_{k=1}^{K}\beta_k\text{Demographic}_{ik}$
$+ \Sigma_{l=1}^{L}\gamma_l\text{VideoInfo}_{il}$

ii. $\text{HighRankingScore}(i) = \alpha_0 + \Sigma_{j=1}^{J}\alpha_j\text{Treatment}_{ij}$
$+ \Sigma_{k=1}^{K}\beta_k\text{Demographic}_{ik}$
$+ \Sigma_{l=1}^{L}\gamma_l\text{VideoInfo}_{il}$

iii. $\text{Satisfaction}(i) = \alpha_0 + \Sigma_{j=1}^{J}\alpha_j\text{Treatment}_{ij}$
$+ \Sigma_{k=1}^{K}\beta_k\text{Demographic}_{ik}$
$+ \Sigma_{l=1}^{L}\gamma_l\text{VideoInfo}_{il}$

As we show in Table 3, the effect of the dual-network treatment is consistently significant and positive on all measures. The user-generated-links-only condition (C1.3) has no significant effect on any of these indexes, which further supports our hypothesis.

In addition, we analyzed the average number of videos participants watched in each condition. Our results show that participants who used the dual-network condition tended to watch fewer videos on average compared with participants in the other conditions (9 videos compared with 10.4 videos, respectively; $p < .01$). Because participants were not required to watch each video in its entirety and were able to move on to subsequent videos before reaching the end of the previous video, this may indicate that dual-network participants liked the videos presented to them and thus spent more time viewing each video. Therefore, we extended the analysis by investigating the average amount of time each participant spent watching each video. We found that, as expected, participants of the dual-network condition spent more time watching each video than participants in the other conditions (126.7 seconds compared with 109.6 seconds, respectively; $p < .01$). This significant observation further supports our findings that the dual network not only helps users find interesting content faster but also serves more interesting content throughout the process.

Figure 6
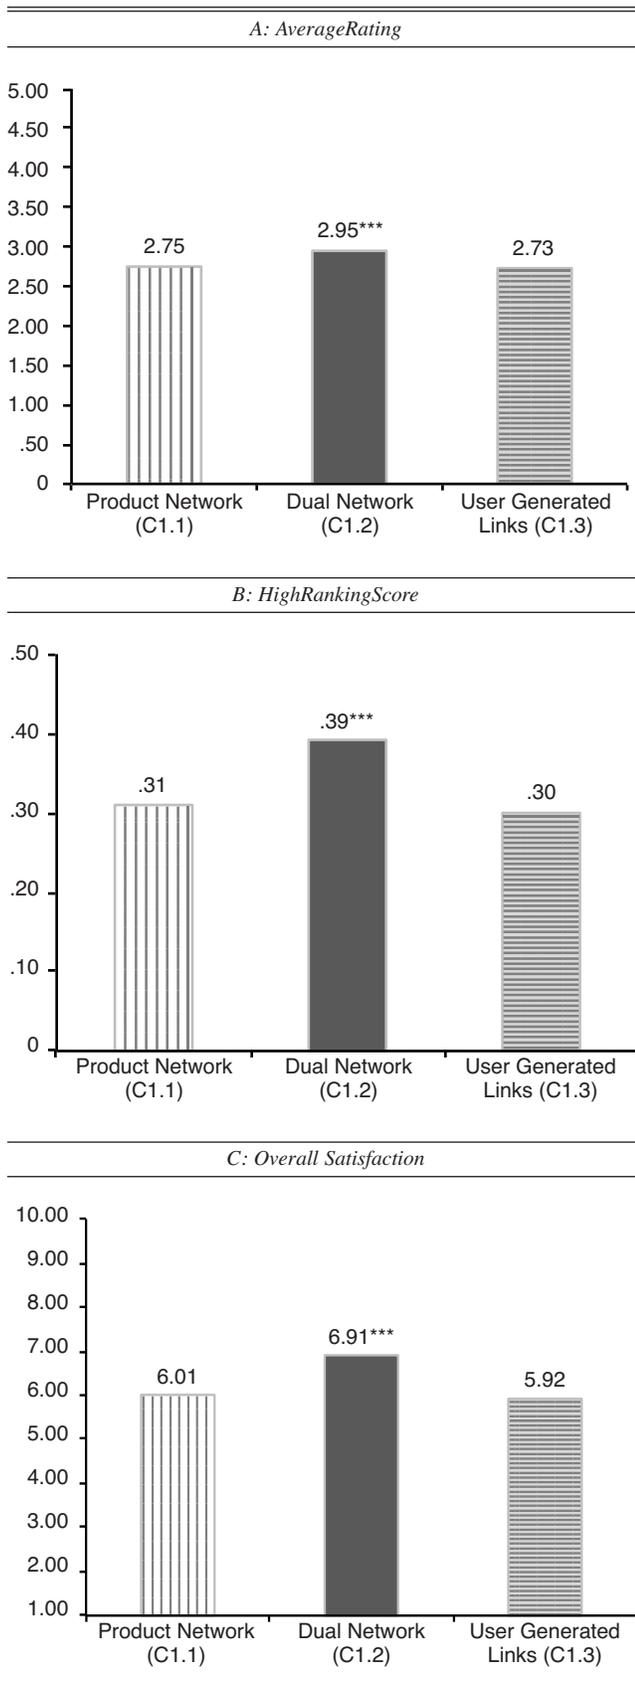STUDY 1: COMPARISON OF PARTICIPANTS' OVERALL
EXPLORATION EFFECTIVENESS

A: AverageRating

B: HighRankingScore

C: Overall Satisfaction

Table 3
STUDY 1: OVERALL EXPLORATION EFFECTIVENESS:
ORDINARY LEAST SQUARES REGRESSION RESULTS

	Overall Satisfaction B (SE)		AverageRating B (SE)		High-RankingScore B (SE)	
Constant	.63	(2.42)	1.107	(.75)	−.35	(.24)
Condition						
Dual network (C1.2)	.73***	(.22)	.18***	(.09)	.06***	(.02)
Only user-generated links (C1.3)	−.23	(.29)	−.04	(.09)	−.02	(.03)
Demographics						
Gender	.11	(.2)	−.08	(.06)	−.03	(.02)
Age	.04	(.02)	.0	(.0)	.0	(.0)
Familiar with YouTube	.36***	(.06)	.07***	(.02)	.02***	(.01)
VideoInfo						
Start from popular video	.28	(.2)	.12*	(.06)	.03	(.02)
First video rating on YouTube	.5	(.5)	.25	(.16)	.12**	(.05)

*Significance at .1 level.
**Significance at .05 level.
***Significance at .01 level.

The Effect on Exploration Stickiness

As we mentioned previously, it may be in the website's best interest to engage the user with the website for a longer period of time. Therefore, we also investigated the effect of the different conditions on the propensity of users to continue using the website. First, we asked participants about their willingness to continue the use of the website (on a ten-point scale). We found that participants who used the dual-network condition ranked their willingness to continue significantly higher (6.81) than participants in the other two conditions (6.31 in the product-network condition and 6.34 in the user-generated-links condition; $p < .05$).

In addition, we subsequently altered the design of the website, allowing users to stop their exploration at any point during the experiment and quit. We used this website design in two settings: In one, we offered the same payment of $1.50 regardless of the time spent on the website, and in the other, the payment to the participant was composed of a basic payment of $.50 and additional payment as a function of the time spent on the website ($.10 per minute up to $1.50). The results of this study show that participants of the dual network are significantly less likely to use the option to leave the website (15% compared with 35% in the product network condition; $p < .01$). Notably, even among those who chose to leave the website, dual-network participants spent more time on the website before choosing to leave it (10.3 minutes compared with 7.6 minutes in the product-network condition; $p < .01$). Taken together, these results indicate that users who used the dual network for content exploration were likely to find content that matched their preferences more quickly and that they produced overall greater satisfaction, and those users reported stronger willingness to continue the exploration process.

To ensure that this was a broad phenomenon, not confined only to the setting and subjects of Amazon's Mechanical Turk, we replicated the same study in a laboratory setting. We randomly assigned 351 undergraduate and MBA students (mean age = 24.5 years, 51.5% female) to the same conditions we described previously. We performed the same

analyses, which yielded essentially the same results, increasing the validity of the general proposition that user-generated links add value to the exploration process.

Study 1a: Replicating the Experiment Using Last.fm Data

To test whether the effect found in Study 1 is unique to YouTube, we replicated the experiment using data from the Last.fm website. Last.fm is a leading music website that offers song recommendations to its more than 40 million users. In the Last.fm product network, each node is a song, and these nodes are linked by the website's recommendation system (links are presented under the title "Similar Tracks"). Registered users in Last.fm are able to create personal webpages and to connect their web pages to those of other users. Users can also generate links to songs. This structure is a dual network, as explained previously; thus, it is suitable for testing our hypothesis and verifying the robustness of our findings outside YouTube. Moreover, because the Last.fm website focuses on music and offers mainly proprietary content and not UGC, it provides an opportunity to test our conjecture outside the UGC domain.

We collected data for a random sample of 50,000 users, 150,000 song tracks, and more than 1 million links between them from the Last.fm website. Almost every track on Last.fm has an embedded link to a YouTube video of that song. We replicated the experiment described in Study 1, using the Last.fm dual network and the Last.fm product network. We asked 120 respondents (47% female, mean age = 29.7 years) to watch music clips for approximately 15 minutes and assign a similar rating ranging from one star ("Poor") to five stars ("Awesome!") to each clip. We randomly assigned each participant to one of two possible conditions: C1a.1 (the Last.fm product network) and C1a.2 (the Last.fm dual network). Note that to be able to use the same experimental design, we presented participants with the video version of each song and not the audio version. We took the recommendations (links) of both the dual network and the product network from the Last.fm website and made no changes to the Last.fm network structure.

We repeated the analysis of Study 1. First, we performed a survival analysis using the Cox proportional hazard model, estimating the effect of each condition on the hazard function. We found that the hazard ratio for the dual-network condition was consistently higher than that of the product network (1.527, $p < .05$). Second, investigating the effect on the overall exploration effectiveness, we observed that the average rating of participants of the dual network was significantly higher than the average of participants of the product network (2.94 for the dual network compared with 2.61 for the product network; $p < .05$). Furthermore, the overall satisfaction of the dual-network participants was significantly higher than that of the product-network participants (7.17 for the dual network compared with 6.42 for the product network; $p < .05$).

These findings refute the possibility that the incremental benefit of the dual network is a result of the algorithm underlying the product network, because YouTube and Last.fm use different algorithms to generate recommendations. Moreover, these results further support our previous findings, extending them to the context of proprietary music.

Study 2: Control for Quality: Testing Whether Effects Are Solely a Result of Added Variety

An alternative explanation for the superior performance of the dual network is that it simply adds variety to the recommendations of the product network and the links have no meaningful contribution to the search. This variety might stem from two sources. First, the dual network is basically two recommendation systems: that of the product network and that of the user-generated links. Second, the user-generated links are not aimed at linking the most similar products (as shown in the level of assortative mixing mentioned previously), and thus the dual network may inherently include more variety. Therefore, it is possible that this added variety caused the effect found in Study 1. In Study 2, we created two dual-network structures and compared them with the YouTube dual network to address this alternative explanation.

Method. Similar to Study 1, we asked participants to watch videos for approximately 15 minutes and assign a rating to each video according to the following new conditions:

- Last.fm links: In this condition, 100 participants (48% female, mean age = 31.5 years) were offered recommendations based on the YouTube product network in addition to recommendations from the Last.fm product network (C2.1). To provide these recommendations, we complemented the data collected from Last.fm by collecting data about the YouTube product-network links for each music video in the Last.fm data set. (The data collection process resulted in 500,000 links from YouTube's product network. For each Last.fm song, we included videos up to two links away using the same data collection method described previously.) We compared this condition with the Last.Fm dual-network condition (C1a.2). The rationale behind this condition is that if the sole contribution of the dual-network is the serving of two types of recommendations, the new condition might be able to provide the same benefits as the dual network.
- Randomized links: In this condition, 300 participants (51% female; mean age = 26.5 years) were offered recommendations based on the YouTube product network in addition to randomly suggested videos (C2.2). That is, under the label "Related Videos," they were offered five recommendations based on the YouTube product network, and under the label "Featured Videos," they were offered five recommendations of randomly selected videos (see Figure 4, Panel D). Here, we compared the YouTube dual-network condition (C1.2) of Study 1. The rationale behind this condition is that if the sole contribution of the dual network is the diversity it offers, the new condition might be able to provide the same benefits as the dual network.

Analysis and results. For each condition, we followed the analysis of Study 1 and computed the hazard ratio. In both cases, as in Study 1, the hazard ratio for the dual-network condition was higher than that of the other conditions (hazard rate of 1.773 compared with the Last.fm links condition and 1.201 compared with the randomized links condition; $p < .05$). We also found that overall satisfaction was significantly higher among participants in the dual-network condition (C1.2) than among participants in the Last.fm-links condition (C2.1) (1.1 difference, $p < .05$) and those in the randomized-links condition (C2.2) (1.0 difference, $p < .05$). Similarly, the AverageRating and the HighRankingScore were higher for the dual-network condition, though the differences were only marginally significant (in the Last.fm-

links condition, the AverageRating was significantly lower, with a difference of .27; $p < .1$).

These results provide additional evidence that merely adding variety, either through a complementary recommendation algorithm or through a randomized network, cannot mimic the benefits of the real-world dual network. A possible explanation for our result is that while random links add variety, they also increase search costs by adding less interesting options. User-generated links are better able to add variety while imposing lower costs on the consumer.

Study 3: Control for Quality: Testing Whether Effects Are Solely a Result of Higher Quality

Another alternative explanation for the superior performance of the dual network might stem from the better quality of the content recommended by users. It is reasonable to expect some correlation between the quality of a content item and the likelihood of that item being labeled as a favorite by a person. Therefore, when a user explores the dual network, she browses through content that has been prefiltered for quality, creating an opportunity to find "good" content sooner. We designed Study 3 to address this alternative explanation.

Method. Similar to Study 1, 100 participants (45% female, mean age = 25.7 years) watched videos for approximately 15 minutes and assigned a rating to each one. Here, we compared the dual-network (C1.2) conditions of Study 1 with a new condition: the high-rating-product-network condition (C3.1). In this condition, participants were offered recommendations based on the product network, in which the suggested videos included only videos rated as above average. All recommendations were shown under the label "Related Videos."

Analysis and results. We followed the analysis of Study 1 and computed the hazard ratios of these two conditions. The hazard for the dual-network condition was significantly higher than that of the high-rating-product-network condition (1.278, $p < .05$). Here, again, overall satisfaction was significantly higher among participants in the dual-network condition (C1.2) than among those in the high-rating-product-network condition (C3.1) (6.9 and 6.0, respectively, $p < .05$). Similarly, the AverageRating was also significantly higher among participants of the dual-network condition (2.95 compared with 2.76 in the C3.1 condition). The HighRankingScore was also higher for the dual-network condition, though the differences were not significant.

Study 4: Testing Whether Effects Are a Result of the Labels Attached to the Content

The remaining question is whether the observed improvement in satisfaction and rating was indeed a result of the specific content recommended by the dual network links or whether it occurred because participants attached some meaning to the knowledge that the links were created by other people and not by an obscure, impersonal algorithm. It might be that labeling primes the user to perceive the content and value of the recommendations in a certain way. For example, a user may feel more emotionally attached to (and thus more influenced by) information that comes from a person then to information generated by a machine.

To test this idea, Study 4 compares both the product-network (C1.1) and the dual-network (C1.2) conditions of Study 1 with five additional conditions. First, we offered the user the dual-network recommendation (similar to C1.2) but changed the label of the user-generated links from "People who liked this also liked" to "Videos liked by friends of this video author" (C4.1). Second, we offered the user the dual-network recommendation (similar to C1.2), but we switched the labeling of the two groups of recommendations: We marked links from the product network as "People who liked this also liked," and we labeled links based on the user-generated links as "Related Videos" (C4.2). Third, we offered the user the dual-network recommendation (similar to C1.2), but all recommendations were provided under the same label ("Related Videos") (C4.3) or with no labels (C4.4). Fifth, participants were exposed to ten recommendations based only on the YouTube product network (similar to C1.1), all under the label of "People who like this also liked" instead of "Related Videos" (C4.5). Thus, the first four conditions, C4.1–C4.4, offered the same recommendations as the dual-network condition (C1.2) described in Study 1 but disguised under different labels. The fifth condition (C4.5) offered the same recommendations as the product-network condition described in Study 1 (C1.1).

We carried out the study according to the procedure described in Study 1, with 500 participants (48.5% female, mean age = 28.3 years), each assigned randomly to one of the new conditions. Using the same criteria as in Study 1, we removed 7 participants from the data set. As in Study 1, we compared the hazard ratios of the different conditions. The results show that participants who were exposed to dual-network recommendations (C1.2, C4.1, C4.2, C4.3, and C4.4) found "good" content more quickly than did those who were exposed to the product-network recommendations only (C1.1 and C4.5), no matter how those recommendations were labeled (hazard ratios for the dual-network conditions: $HR_{C1.2} = 1.285$, $HR_{C4.1} = 1.249$, $HR_{C4.2} = 1.217$, $HR_{C4.3} = 1.296$, and $HR_{C4.4} = 1.38$; the hazard ratio of the product network labeled as the dual network is $HR_{C4.5} = .957$; all are compared with the product-network condition $HR_{C1.1}$ as a baseline; $p < .05$).

We also performed a contrast analysis (Rosenthal and Rosnow 1985), which revealed a significant difference (F = 5.98, $p < .01$) in overall satisfaction in favor of participants who used the dual network (C1.2 and C4.1–C4.4) compared with participants exposed to the product-network structure only (C1.1 and C4.5). In addition, we found significant differences in the AverageRating and the HighRankingScore variables (F = 4.26, $p < .01$, and F = 3.55, $p < .01$, respectively) when contrasting participants in the dual-network condition with those in the product network condition. There were no significant differences for those measurements for participants in the different dual-network conditions (C1.2 and C4.1–C4.4) and no significant difference among participants in the product-network conditions (C1.3 and C4.5). These findings show that participants exposed to dual-network recommendations performed better in finding "good" content than did those who were exposed to the product-network recommendations, regardless of the labeling of the recommendations, implying that the mere labeling has a marginal effect, if any.

THE EFFECT OF DYNAMIC STRUCTURE CHANGES

Unlike stores, electronic commerce sites (e.g., Amazon. com) can respond to consumers' searches in real time and change displays dynamically to direct their exploration. Furthermore, as Hui, Bradlow, and Fader (2009) observe, as the amount of time a consumer spends in a store increases, the consumer's tendency to continue exploring the store decreases. A possible solution to this problem is to dynamically change the "store" design throughout the consumption process. Thus, rather than present consumers with a static view of the product network, e-commerce sites can dynamically incorporate user-generated links to improve the efficiency of the search process and to retain customers' interest. When such a dynamic strategy is used, the question becomes this: Is there a point during the exploration process at which the consumer might benefit more from the introduction of user-generated links?

We designed the following studies to provide initial answers to this question. In these studies, each participant was exposed to different structures at different times during the navigation process. The switch to a new network structure was based on either the phase of the exploration process the user was in (Studies 5 and 6) or his or her satisfaction during exploration (Study 7).

Study 5: Structure Change as a Function of Exploration Time

In Study 5, we explore the conjecture that the dual-network structure will be most beneficial to consumers who have been browsing the website for a while and may need to revive the search with more diversity. This conjecture is based on the belief that in a continuous process of exploration, product network recommendation mechanisms tend to narrow down the navigation possibilities to objects that are similar to one another. This is also supported by our results with respect to the high level of assortative mixing mentioned previously. Therefore, we created a new condition (C5.1) using a dynamic dual network and compared it with the previously mentioned static conditions (product network, dual network). Under this new condition, consumers were presented with recommendations based on the product network structure and, after a specified amount of time, were switched to the dual-network structure.

Method. In this study, we assigned 300 participants (49% female; mean age = 28.5 years) randomly to one of the following conditions. One-third of the participants were presented with the dynamic dual-network condition (C5.1), with the switch from product-network to dual-network recommendations occurring after 250 seconds. The rest of the participants were presented with the static product-network (C1.1) and static dual-network (C1.2) conditions. The experimental procedure was the same as that described in Study 1. Using the same criteria as in Study 1, we removed 11 participants from the data set. As in Study 4, to control for a possible labeling effect and to avoid a visible change in the appearance of the website in the middle of the study, we grouped all types of recommendations under the same label ("Related Videos").

Analysis and results. Given that this study was a within-subject design, the hazard model estimated in the previous section was no longer appropriate. In addition, we could not simply compare the ratings before and after the change in network structure, because our previous results demonstrated a continuous decrease in users' rankings over time regardless of the types of recommendations presented. Thus, to replicate the previous results (dual-network effect) we needed to compare between and across subjects, using a difference-in-differences model (Bertrand, Duflo, and Mullainathan 2004). In our context, the control group was the group presented with the product-network recommendations only, with no dynamic switch (i.e., C1.1). We estimated the following model:

$$\text{Rating}_{i,t} = \alpha + \beta \text{Treatment}_i + \gamma \text{Period}_t$$
$$+ \delta \text{Treatment}_i \times \text{Period}_t + \varepsilon_{i,t},$$

where $\text{Rating}_{i,t}$ is the average rating in period t for participant i; Treatment_i is a treatment dummy variable, with $T_i = 1$ representing exposure to the dynamic condition; Period_t is a period dummy variable, where $\text{Period}_t = 0$ indicates the time period in the study before treatment (i.e., the time period before the switch to the dual network; $t \le 250$ seconds) and $\text{Period}_t = 1$ indicates the time period after treatment ($t > 250$ seconds); and $\text{Treatment}_i \times \text{Period}_t$ is an interaction term and represents the actual treatment variable. Table 4 presents the estimates of this regression.

The coefficient for Treatment is not significant, indicating that there are no significant differences between the groups (treated and untreated) a priori. Similarly, the coefficient for Period is not significant, indicating that there are no significant differences within subjects who are not treated. The interaction coefficient is positive and significant, thus indicating that users under the dynamic condition demonstrated a significant increase in rating after the switch to the dual network.

As Table 5 shows, users of the dynamic dual-network design also had higher satisfaction than did users of the

Table 4

STUDY 4: DIFFERENCE-IN-DIFFERENCES ESTIMATION OF USERS' RATINGS UNDER THE DYNAMIC DUAL-NETWORK CONDITION (C5.1)

Parameter	Estimate	SE	Significance
(Intercept)	2.93	.06	.00
Treatment	.09	.13	.49
Period	−.09	.09	.28
Period × Treatment	.42	.19	.02

*Significance at the .05 level.

Notes: F = 4.909, significance < .01.

Table 5

STUDIES 5–6: COMPARISON OF THE OVERALL SATISFACTION, AVERAGE RATING, AND HIGH RANKING SCORE OF PARTICIPANTS OF THE DYNAMIC BY TIME CONDITIONS

	Overall Satisfaction	AverageRating	HighRankingScore
Product network (C1.1)	6.0	2.74	.31
Dual network (C1.2)	6.9	2.95	.39
Time-dependent dynamic dual (C5.1)	7.5	2.9	.37
Time-dependent dynamic randomized links (C6.1)	6.3	2.71	.32

Notes: All estimations use post hoc ANOVA using Fisher's least significant difference test.

static product-network design ($p < .01$) or the static dual-network design ($p < .01$). Observing the AverageRating and the HighRankingScore, which are measures that represent the entire process, we find that participants in condition C5.1 scored significantly higher than did those in the static product network condition ($p < .05$). However, we find no significant difference between participants of this time-dependent dynamic dual-network condition (C5.1) and those of the static dual network condition (C1.2).

Study 6: Testing Whether the Effect Is Driven by the Content of the Generated Links or Whether This Is an Artificial Effect Caused by Mere Change

It is possible that the effects observed in Study 5 are merely the results of diversity: Users may simply enjoy the opportunity to view videos that seem different from the videos previously offered. Thus, the results are not necessarily related to the content of the links in the dynamic condition. Therefore, for robustness, and to rule out this alternative explanation, we carried out Study 6, creating a new dynamic structure based on the randomized links.

Method. Under this condition (C6.1), 100 participants (50% female, mean age = 27.6 years) were presented with recommendations based on the product-network structure and after a fixed amount of time were switched to the randomized-links structure. (We excluded 2 participants on the basis of the criteria described previously.) This is an important control, because if consumers only seek diversity, a randomized-links condition might be able to provide the same benefits as the dual network. Again, to control for a possible labeling effect and to avoid a visible change in the appearance of the website in the middle of the study, we grouped all recommendations under the same label ("Related Videos").

Analysis and results. We estimated a difference-in-differences model similar to the one in Study 5, with the treatment now being the switch to the randomized links. Table 6 presents the results and shows that in contrast to the time-dependent dynamic dual-network condition (C5.1), switching from the product-network to the randomized-links condition had a significant, negative effect, indicating again that the randomized links cannot mimic the positive effect of the user-generated links and the dual-network role. We also compared participants' overall satisfaction, Average-Rating, and HighRankingScore (see Table 5) and found that these measures were significantly lower among participants in the dynamic randomized-links condition (C6.1) than among those in C5.1, the time-dependent dynamic dual-network

Table 6
STUDY 6: DIFFERENCE-IN-DIFFERENCES ESTIMATION OF USERS' RATINGS UNDER THE DYNAMIC RANDOMIZED-LINKS CONDITION (C6.1)

Parameter	Estimate	SE	Significance
(Intercept)	2.93	.07	.00
Treatment	.22	.14	.12
Period	−.1	.09	.30
Period × Treatment	−.39*	.20	.05

*Significance at the .1 level.
Notes: F = 2.984, significance < .05.

condition ($p < .01$), or in C1.2, the static dual-network condition ($p < .05$).

Study 7: Structure Change as a Response to Participants' Discontentedness

Our findings in Study 6 show that switching from a product network to a dual network after a certain period of time in the exploration process has a positive effect. In Study 7, we focus on a dynamic environment in which the site structure changes in response to the consumer's reactions. Like Study 5, this approach is based on the belief that during a continuous process of exploration, product-network recommendation mechanisms tend to narrow down the navigation possibilities to very similar objects. When a user dislikes what he or she is watching, offering similar objects will not be of great help, and it may therefore be time to point the user to different areas in the product network. Thus, introducing a dual network when the user is dissatisfied might improve the exploration process.

Method. In Study 7, the dynamic structure was designed as follows: When a user began the exploration process, the recommendations he or she received were based on a predefined network structure. When (and if) the user ranked two consecutive videos with a low rating (< 3), the structure of the site changed so that the user was exposed to a different network structure.

We randomly assigned 300 participants (51% female, mean age 28.9 years) to the following dynamic conditions: (1) beginning from the product-network and switching to the dual-network structure for the rest of the exploration time (C7.1); (2) beginning from the product network and switching to the randomized links together with the product-network structure for the rest of the exploration time (C7.2); and (3) beginning from the dual-network and switching to the product-network structure for the rest of the exploration time (C7.3). Because the change of structure was induced by the participant's activities, it was possible that a participant would not be exposed to a change in structure at all. We removed 9 participants using the same criteria as in Study 1.

Analysis and results. We found that nearly 80% of the participants experienced a sequence of videos that did not match their preferences and were therefore exposed to a change of structure during their exploration. Because the time of the switch varied across participants, we could not use the difference-in-differences model estimated previously to analyze the results of this study and were therefore unable to control for potential across-subject differences. Nevertheless, we were able to compare within subjects using paired sample t-tests. For each participant, we compared the average ranking before the switch with the average ranking after the switch. As Table 7 shows, participants in condition C7.1 (dynamic dual network by ranking; i.e., shift from product to dual network structure) experienced the largest jump in average rating. Among these participants, the switch increased video ratings by 17.5% on average. In conditions C7.2 and C7.3, we did not find a significant difference between consumers' average ratings before versus after the switch. These findings suggest that a change from the product network to the dual network as a response to the user's indication of dissatisfaction may lead to an increase of his or her ranking of the products he or she con-

Table 7

STUDY 7: THE DIFFERENCE IN AVERAGE RATING BEFORE
AND AFTER THE STRUCTURAL SWITCH

	M	SD	Significance (Two-Tailed)
Dynamic dual by ranking (C7.1)	.42*	.98	.00
Dynamic randomized links by ranking (C7.2)	.11	.84	.30
Dynamic dual to product (C7.3)	−.08	.87	.52

*Significance at the .01 level.

Notes: In condition C7.1, the algorithm switches from product network to dual network; in condition C7.2, the algorithm switches from product network to random network; and in condition C7.3, the algorithm switches from dual network to product network.

sumes. We may further conclude that the increase in the rating is not due to the effect of switching between different structures but is influenced by the actual structure presented to the participant.

Table 8 presents the results with regard to the Average-Rating, HighRankingScore, and overall satisfaction. Given that participants in two conditions (i.e., C7.1 and C7.3) were exposed to the dual network, it was not clear a priori which condition would lead to better results. Therefore, we used a post hoc ANOVA to compare these measures. We found that the average rating of participants in condition C7.1 (i.e., the dynamic dual network by ranking) was significantly higher than the average ranking under any other condition tested in our study ($p < .05$), including the static dual-network condition tested in C1.2 and the time-dependent dynamic dual-network condition tested in C5.1. Similarly, we found that both the overall satisfaction and the HighRankingScore of participants in condition C7.1 were significantly higher than those in any of the other tested conditions ($p < .01$ when compared with C1.1; $p < .05$ when compared with C7.2 and C7.3; and $p < .1$ when compared with C1.2 and C5.1). Moreover, the overall satisfaction was significantly higher ($p < .01$) than in the static dual-network condition (C1.2) and even higher ($p < .1$) than in the time-dependent dynamic dual-network condition (C5.1).

Notably, we found that the overall satisfaction of participants in condition C7.3 was significantly higher ($p < .01$) than the overall satisfaction of participants in the static product-network condition (C1.1) but not significantly different from that in the static dual-network condition (C1.2).

Table 8

STUDY 7: COMPARISON OF THE OVERALL SATISFACTION,
AVERAGERATING, AND HIGHRANKINGSCORE OF
PARTICIPANTS OF THE DYNAMIC BY TIME CONDITIONS

	Overall Satisfaction	AverageRating	HighRanking Score
Product network (C1.1)	6.01	2.72	.32
Dual network (C1.2)	6.93	2.97	.39
Dynamic dual by ranking (C7.1)	8.00	3.14	.42
Dynamic randomized links by ranking (C7.2)	6.34	2.71	.32
Dynamic dual to product (C7.3)	6.88	2.82	.34

Notes: All estimations used Fisher's least significant difference test.

In accordance with our findings in the preceding studies, the overall satisfaction of participants in the ranking-dependent dynamic randomized-links condition (C7.2) was significantly lower than the satisfaction rate in any of the dual-network conditions ($p < .01$ when compared with C5.1 and C7.1; $p < .05$ when compared with C1.2 and C7.3).

DISCUSSION

Product search and exploration have become complicated tasks due to the vast variety and quantity of products available online. The presence of a linked product network, which is one of the principal differences between online and traditional content environments, facilitates consumers' exploration process. We study how this process is affected by the integration of user-generated links into the product network structure, which creates a dual network. We find that user-generated links improve exploration efficiency by leading consumers to find better content more quickly and improve exploration effectiveness by increasing overall consumer satisfaction. In addition, we propose a dynamic approach to further improve this effect.

A limitation of this study is that our website does not have the keyword search option that is available on YouTube and on most other websites. Because our analysis focuses on exploration, we believe that this is not a major limitation, but further research could study the interaction between directed search and exploration by adding such a feature. Nevertheless, we propose that any recommendation mechanism should still include individual-level data (i.e., user-generated links) and not solely aggregate information, as is typically found in recommendation system algorithms.

We did not find any evidence to indicate why YouTube decided to offer users the opportunity to connect to one another in a social network. We can assume that their goal was increasing views and involvement of users. However, our results suggest that the user-generated links in such networks can improve the support the product network provides to exploration search. Thus, by incorporating social network features, an e-commerce site can produce a static dual-recommendation network with relatively little effort. Alternatively, such sites might adopt a more sophisticated recommendation system that exposes users to user-generated links at different times. Studies 5 and 7 offer indications for managers regarding how and when recommendations from the dual network structure should be integrated into each product page to increase user satisfaction. Such transitions might be time dependent in situations in which the user's level of satisfaction is not known, or they might be satisfaction dependent in cases in which the user expresses his or her opinion in real time regarding the content he or she consumes. Although it does not currently seem feasible to track an individual user over a large span of time (e.g., several months), firms may study and optimize strategies for dynamically offering user-generated links on the basis of observed satisfaction. In cases in which satisfaction is not observable, a dynamic approach based on exploration time can also be used.

A potential caveat regarding any recommendation strategy aimed at increasing search efficiency is that users might sour on the site when they find good content. Given that many content websites use an advertisement-based business model, rapid conversion might not be ideal, and the optimal

strategy might be to keep the user on the site. It could be inferred that if this is correct, a less efficient recommendation network might actually be preferable to an efficient one. However, we show that using the dual-network structure increases not only the efficiency (time to good content) but also the effectiveness (total rating and overall satisfaction). Indeed, users with higher overall satisfaction may spend more time on the website. We explored this possibility in another study (see Appendix B) in which participants were allowed to leave the website at any point during the experiment. We tested two settings: In one, we offered the same payment regardless of the time spent on the website, and in the other, the payment to the participant was composed of a basic payment and additional payment as a function of the time spent on the website. The results of this study show that participants using the dual network were less likely to use the option to leave. In addition, among those who chose to leave the website, participants of the dual network spent more time on the website. However, it is still not clear that a better recommendation system will always lead to a greater amount of time spent on the site. It is also not clear how satisfaction and time spent in a current search influences a future search or even initiating a search. From an indirect perspective, an improved exploration process might prompt more frequent future visits to a website, which might be especially relevant in cases of competition between various e-commerce sites. The effect of user-generated links on repeat visits is beyond the scope of this work. The connection between satisfaction and time spent browsing is indeed a complex phenomenon that should be explored in further research.

APPENDIX A: NETWORK INDEXES DEFINITION

Using the Complex Networks Analysis Package (Muchnik et al. 2007), we computed the following indexes:

1. The degree (number of links) of each node, including the in-degree (the number of incoming links) and out-degree (the number of outgoing links).
2. Closeness centrality of each node: A measure of the average minimal distance (numbers of hops) between this node and any other node in the network.
3. Betweenness centrality of each node: A measure of the number of shortest paths (between any two nodes) in which this node is included. The formal definition of the betweenness (BC) of node i is as follows:

$$BC_i = \sum_{s < t} \frac{g_i^{(st)}}{n_{st}},$$

where n_{st} is the number of shortest paths from node s to node t and $g_i^{(st)}$ number of shortest paths from s to t that pass through node i.

4. PageRank of each node (Brin and Page 1998): An iterative measure of centrality, which is based on the number of links pointing to a node and their centrality. The formal definition of the PageRank of node i is as follows:

$$PR(i) = \sum_{v \in B_i} \frac{PR(v)}{L(v)},$$

where v is a node in the set of all nodes linking to node i (denoted as B_i). PR(v) is the PageRank of node v, and L(v) is all outgoing links of node v.

5. Assortative mixing level of each node: A measure of the level of similarity (homophily) between a given node and its neighbors (Newman 2003). We operationalize a product's assortative mixing as the percentage of its neighboring products that are of the same category (using the 14 categories defined by YouTube).

APPENDIX B: STUDY 8: ALLOWING QUITTING THROUGHOUT THE EXPLORATION

We performed a supplemental study replicating the design of Study 1 but allowing participants to leave during the experiment. We designed this study to provide insight into consumer behavior in a setting that is more like a real website, in which efficiency and satisfaction influence how long the consumer spends browsing, as opposed to a setting in which all participants are required to spend a fixed amount of time browsing the site.

Method

In this study, we assigned 200 participants randomly to two payment settings with two conditions each, the product network and the dual network, using Study 1's design. In both settings, after five minutes of exploration, participants were able to leave the experiment and be rewarded for their participation. In the first design, all participants were paid $1.50 regardless of the time spent on the site. In the second design, participants received a basic payment of $.50 and an additional $.10 for every additional minute spent on the site.

Analysis and Results

We estimated the probability of leaving the site in each of the conditions, comparing the product-network conditions (C8.1) with the dual-network conditions (C8.2). The results show that in both payment settings, participants who used the dual-network conditions were significantly less likely to leave the website during the exploration process (15% compared with 35% in the product network, $p < .01$). For those participants who chose to leave, we also compared the time spent on the site before choosing to leave. We found that dual-network participants who left the website spent significantly more time on the website than did product-network participants. Specifically, among participants who chose to leave, dual-network participants spent an additional 5.3 minutes on the site after the option to leave was offered to them, for a total of 10.3 minutes. Product-network participants, in contrast, spent only 2.6 minutes in addition to the mandatory 5-minute period ($p < .01$). These results show that even when offered the possibility to leave the website during exploration, a situation that is closer to a real-world website, dual-network users were more satisfied, as reflected by their lower tendency to leave the website as well as the total time spent on the site

REFERENCES

Adomavicius, Gediminas and Alexander Tuzhilin (2005), "Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art and Possible Extensions," *IEEE Transactions on Knowledge and Data Engineering*, 17 (6), 734–49.

Ahn, Yong-Yeol, Seungyeop Han, Haewoon Kwak, Sue Moon, and Hawoong Jeong (2007), "Analysis of Topological Characteristics of Huge Online Social Networking Services," in *Proceedings of the 16th International Conference on World Wide Web*.

New York: Association for Computing Machinery, 835–44.

Arasu, Arvind, Junghoo Cho, Hector Garcia-Molina, Andreas Paepcke, and Sriram Raghavan (2001), "Searching the Web," *ACM Transactions on Internet Technology*, 1 (1), 2–43.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan (2004), "How Much Should We Trust Differences-in-Differences Estimates?" *Quarterly Journal of Economics*, 119 (1), 249–79.

Brin, Sergey and Lawrence Page (1998), "The Anatomy of a Large-Scale Hypertextual Web Search Engine," *Computer Networks ISDN*, 30 (1–7), 107–117.

Bucklin, Randolph E., James M. Lattin, Asim Ansari, Sunil Gupta, David Bell, Eloise Coupey, et al. (2002), "Choice and the Internet: From Clickstream to Research Stream," *Marketing Letters*, 13 (3), 245–58.

Buldyrev, Sergey V., Roni Parshani, Gerald Paul, H. Eugene Stanley, and Shlomo Havlin (2010), "Catastrophic Cascade of Failures in Interdependent Networks," *Nature*, 464 (7291), 1025–28.

Burt, Ronald S. (1992), *Structural Holes: The Social Structure of Competition*. Cambridge, MA: Harvard University Press.

Carrington, Peter, John Scott, and Stanley Wasserman (2005), *Models and Methods in Social Network Analysis*. New York: Cambridge University Press.

Cox, D. (1972), "Regression Models and Life-Tables," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 34 (2), 187–220.

Freeman, Linton C. (1977), "A Set of Measures of Centrality Based upon Betweenness," *Sociometry*, 40 (1), 35–41.

Goldenberg, Jacob, Sangman Han, Donald R. Lehmann, and Jae Weon Hong (2009), "The Role of Hubs in the Adoption Process," *Journal of Marketing*, 73 (March), 1–13.

—— , Barak Libai, and Eitan Muller (2001), "Talk of the Network: A Complex Systems Look at the Underlying Process of Word-of-Mouth," *Marketing Letters*, 12 (3), 211–23.

Granovetter, Mark (1973), "The Strength of Weak Ties," *American Journal of Sociology*, 78 (6), 1360–80.

Herlocker, Jonathan L., Joseph A. Konstan, Loren G. Terveen, and John T. Riedl (2004), "Evaluating Collaborative Filtering Recommender Systems," *ACM Transactions on Information Systems*, 22 (1), 5–53.

Hui, Sam K., Eric T. Bradlow, and Peter S. Fader (2009), "Testing Behavioral Hypotheses Using an Integrated Model of Grocery Store Shopping Path and Purchase Behavior," *Journal of Consumer Research*, 36 (3), 478–93.

—— , Peter S. Fader, and Eric T. Bradlow (2009), "Path Data in Marketing: An Integrative Framework and Prospectus for Model Building," *Marketing Science*, 28 (2), 320–35.

Kaplan, E.L. and Paul Meier (1958), "Nonparametric Estimation from Incomplete Observations," *Journal of the American Statistical Association*, 53 (282), 457–81.

Katona, Zsolt and Miklos Sarvary (2008), "Network Formation and the Structure of the Commercial World Wide Web," *Marketing Science*, 27 (5), 764–78.

—— , Peter Pal Zubcsek, and Miklos Sarvary (2011), "Network Effects and Personal Influences: Diffusion of an Online Social Network," *Journal of Marketing Research*, 48 (June) 425–43.

Kleinbaum, David G. and Mitchel Klein (2005), *Survival Analysis: A Self-Learning Text*, 2d ed. New York: Springer.

Moe, Wendy (2003), "Buying, Searching, or Browsing: Differentiating Between Online Shoppers Using In-Store Navigational Clickstream," *Journal of Consumer Psychology*, 13 (1/2), 29–39.

—— and Peter S. Fader (2004), "Dynamic Conversion Behavior at E-Commerce Sites," *Management Science*, 50 (3), 326–35.

Muchnik, Lev, Royi Itzhack, Sorin Solomon, and Yoram Louzoun (2007), "Self-Emergence of Knowledge Trees: Extraction of the Wikipedia Hierarchies," *Physical Review E*, 76 (1), (published electronically July 13), [DOI: 10.1103/PhysRevE.76.016106].

Newman, M.E.J. (2003), "The Structure and Function of Complex Networks," *SIAM Review*, 45 (2), 167–256.

Oestreicher-Singer, Gal and Arun Sundararajan (2012a), "Recommendation Networks and the Long Tail of Electronic Commerce," *MIS Quarterly*, 36 (1), 65–83.

—— and —— (2012b), "The Visible Hand of Social Networks in Electronic Markets," *Management Science*, forthcoming.

Rangaswamy, Arvind, C. Lee Giles, and Silvija Seres (2009), "A Strategic Perspective on Search Engines: Thought Candies for Practitioners and Researchers," *Journal of Interactive Marketing*, 23 (1), 49–60.

Rosenthal, Robert and Ralph L. Rosnow (1985), *Contrast Analysis: Focused Comparisons in the Analysis of Variance*. New York: Cambridge University Press.

Shaikh, Nazrul, Arvind Rangaswamy, and Anant Balakrishnan (2006), "Modeling the Diffusion of Innovations Using Small-World Networks," working paper, Department of Marketing, Penn State University.

Trusov, Michael, Randolph E. Bucklin, and Koen H. Pauwels (2009), "Effects of Word-of-Mouth Versus Traditional Marketing: Findings from an Internet Social Networking Site," *Journal of Marketing*, 73 (September), 90–102.

Valente, Thomas W. (1996), *Network Models of the Diffusion of Innovations*. New York: Hampton Press.

Van den Bulte, Christophe and Stefan Wuyts (2007), *Social Networks and Marketing*. Cambridge, MA: Marketing Science Institute.

Wasserman, Stanley and Katherine Faust (1994), *Social Network Analysis: Methods and Applications*. New York: Cambridge University Press.