Designing Sustainable Online Support: Examining the Effects of Design Change in 49 Online Health Support Communities

Joshua Introne* 
Department of Media & Information, Michigan State University, East Lansing, MI. E-mail: jintrone@msu.edu

Ingrid Erickson 
School of Information Studies, Syracuse University, Syracuse, NY

Bryan Semaan
School of Information Studies, Syracuse University, Syracuse, NY

Sean Goggins
Electrical Engineering and Computer Science, University of Missouri, Columbia, MO

Online social support communities can significantly improve health outcomes for individuals living with disease. Although they are well studied in the literature, little research examines how sociotechnical design changes influence the sustainability of support communities for different medical conditions. We compare the impact of a single design change on 49 disease-specific health support forums hosted on the WebMD platform, a popular online health information service. A statistical analysis showcases changes in posting patterns before and after the design intervention; a subsequent interpretive examination of forum content reveals how the design change affected members' perceived affordances of the platform. Our findings suggest that, despite differences between communities, the design change triggered a common set of cascading effects: it made it difficult for core users to create and maintain relationships, that led them to ultimately leave the site, and, in turn, reduced the activity drawing newcomers to the platform. Using these findings, we argue that the design of sustainable and robust online communities must account for systemic, sociotechnical dynamics.

*Corresponding author

Additional Supporting Information may be found in the online version of this article.

Received August 20, 2018; revised February 14, 2019; accepted April 14, 2019

 DOI: 10.1002/asi.24250

Introduction

Roughly 150 million U.S. citizens have one or more chronic diseases and over half of these individuals have sought support for their disease on the Internet (Fox & Purcell, 2010). Online participation often begins when individuals navigate to an online health community (OHC) in search of information, but typically evolves to encompass the exchange of social support (Costello, Martin, & Edwards Brinegar, 2017). Social support is a multidimensional construct comprising several different kinds of aid, regard, or assistance that peers provide to one another via social networks (Fox & Purcell, 2010).

The value of social support within the realm of health and wellness has long been recognized (Kaplan, Cassel, & Gore, 1977) and OHCs have organically responded to this need (Taylor, 2011) by becoming places where people with medical conditions go to find social support that complements the services provided by professional health care providers (Erfani, Abedin, & Blount, 2017; Garcia, Mavrodiev, & Schweitzer, 2013; Stvilia, Mon, & Yi, 2009). Today, tens of thousands of online communities have emerged around the exchange of social support for health and wellness (as reported by Eysenbach, Powell, Englesakis, Rizo, & Stern, 2004).

Many types of support may be given in OHCs, but scholars often distinguish between *informational* and *socio-emotional* support (Nakikj & Mamykina, 2018). The former involves the provision of knowledge or advice, and the latter includes direct emotional support and that which comes with the sense of belonging to a community (Bambina, 2007; Nakikj & Mamykina, 2018). Patterns of social support have

been found to vary across different OHCs, (Civan, McDonald, Unruh, & Pratt, 2009; Gill, 2012; Hartzler & Pratt, 2011; Huh, McDonald, Hartzler, & Pratt, 2013; Maloney-Krichmar & Preece, 2005) with different communities exhibiting different mixtures of support types (Introne, Semaan, & Goggins, 2016), degrees of responsiveness (Nambisan, Gustafson, Hawkins, & Pingree, 2016), and levels of information quality (Adams, 2010; Berland et al., 2001; Eysenbach, Powell, Kuss, & Sa, 2002).

Although this evolving, deep understanding of how people are drawn to online health communities and enact support in different ways for specific conditions reflects significant progress, less is known about the sociotechnical features of online communities that leads to their sustainability—that is, the ways in which online community spaces come to exhibit stable patterns of social interaction, and are able to grow and/or maintain their user base relative to community design (Butler, 2001; Garcia et al., 2013). One reason for this is that it is hard to disentangle the relative influence of design from that of the health condition discussed and other emergent social factors when comparing different communities. As pointed out by Nakikj and Mamykina (2018), natural experiments wherein a single community experiences different sociotechnical designs are rare.

In this article, we explore how design can influence the sustainability of an OHC. We examine the case of health support forums for forty-nine specific conditions hosted on a single platform that underwent a major design change. Our work is primarily data-driven, and we use longitudinal postings pattern analysis of members in different communities, interrupted time series analysis, regression modeling, and content analysis, to holistically focus on the question: *How might a design change impact the sustainability of an online health community?* In our data, we find that following the design change, each forum experienced different trajectories of decline, and we disentangle the general impact of the design change from specific aspects of each forum. Through forensic analysis of member interactions before and after the design change, we illustrate its differential effects across types of users and disease communities. Our analysis sheds light on a sociotechnical system some communities held in common and helped to sustain them. Further, the analysis also illuminates how a design intervention can evolve into a kind of sociotechnical wave that spreads within each community, ultimately disrupting the sustainability of the system.

Literature Review

Online Communities as Heterogeneous Sociotechnical Systems

Online communities are made up of individuals who engage in support activities for very different reasons and this diversity is reflected in the ways that they participate and, ultimately, organize themselves. Certain individuals join an OHC in search of information or support which they can "take" from their community experience. In many cases, these individuals are newcomers—a type of community actor that has been studied extensively from contexts as varied as organizational listsersvs (Ahuja & Galvin, 2003), FLOSS development communities (Qureshi & Fang, 2011), usenet forums (Burke, Kraut, & Joyce, 2010), and online participatory platforms, such as Wikipedia (Choi, Alexander, Kraut, & Levine, 2010; Halfaker, Geiger, & Terveen, 2014; Halfaker, Keyes, & Taraborelli, 2013). As mentioned, newcomers most frequently cite the desire to find information (Coulson, 2005; Rodgers & Chen, 2005), obtain emotional support (Wang, Kraut, & Levine, 2012), and form relationships (Shim, Cappella, & Han, 2011) as their primary motivations for participation. In the case of OHCs, obtaining informational support is of particular importance (Chung, 2014; Introne et al., 2016).

When (and if) newcomers remain and mature into more seasoned community members, their experiences begin to shift. Often they begin to form relationships with other members (Arguello et al., 2006; Burke et al., 2010; Burke, Joyce, Kim, Anand, & Kraut, 2007; von Krogh, Spaeth, & Lakhani, 2003) and engage in self-directed and self-actuated learning to become familiar with their respective community's quality standards for participation (Choi et al., 2010; Kraut, Burke, Riedl, & Resnick, 2012). Often, they begin to take on more visible roles within the community at this time. Bambina (2007) found that members in a breast cancer OHC organized into "providers" who repeatedly deliver support to "takers," who tend to repeatedly receive support. In studying a support group for people with knee problems, Maloney-Krichmar and Preece (2005) identified 17 distinct roles played by members, such as "gatekeepers" and "group commentators," both of which focus on community maintenance. This social organization of roles has been extensively documented in other communities, such as open-source software communities (K. Crowston, Wei, Li, & Howison, 2006; Mockus, Fielding, & Herbsleb, 2002), Wikipedia (Gorbatai & Piskorski, 2012), as well as in OHCs (Bambina, 2007; Introne et al., 2016), which might be taken as a characterization of how online sociotechnical systems self-organize (e.g., Ackerman, 2000; Gleave, Welser, Lento, & Smith, 2009; Welser, Gleave, Fisher, & Smith, 2007; Yukawa, 2006).

The enactment of different roles within an online community creates interlocking patterns of activity that gives rise to an observable, mesoscale social structure. Studies of OHCs have identified a variety of such structures (Carron-Arthur, Ali, Cunningham, & Griffiths, 2015). One structure, found in OHCs and many other online communities, is described as a *core-periphery* network structure. A core-periphery network is a pattern in which a subset of individuals form a highly interconnected core that is surrounded by a (generally much larger) set of more peripheral individuals that interact with members of the core, but not one another (Borgatti & Everett, 2000; Csermely, London, Wu, & Uzzi, 2013; Cucuringu, Rombach, Lee, & Porter, 2014). Introne et al. (2016) extended the core-periphery model in the context of WebMD to introduce an even larger, more peripheral group of users called the *extraperiphery* that posted a

DOI: 10.1002/asi

handful of times, but never formed relationships on the site. In Bambina's (Bambina, 2007) breast cancer OHC analysis, support providers form the core and support takers form the periphery. Other research documents similar role-related core-periphery structures (K. Crowston et al., 2006; Fisher, Smith, & Welser, 2006; Gorbatai & Piskorski, 2012; Mockus et al., 2002). Introne et al. (2016) showed that core members engage in lengthy, informal conversations, and provide the bulk of the information support for more peripheral members.

Mechanisms of Sustainability in Online Communities

The sustainability of an online community is frequently examined relative to an individual's ongoing commitment to participate. Research finds that individuals that find and maintain relationships are more committed to stay as active members in online communities. For example, Cheung and Lee (2009), in studying an online community of teachers and educators, show that social value, expressed as social bonds, has a positive impact on an individual's continued commitment to a community. Lampe, Wash, Velasquez, and Ozkaya (2010) reinforce this finding and note that social relationships often stand in contrast to individuals' initial motivations for participation. It is unsurprising then that features for "friending" function well for the expansion of social networking platform memberships, whereas interactions such as "liking" and commenting keep people engaged (Kabadayi & Price, 2014).

In OHCs, the kinds of support sought and obtained may play a role in attachment to a community. Wang et al. (2012) found that exposure to emotional support is correlated with length of membership, corroborating Nakikj and Mamykina (2018) qualitative findings about members' prioritization of socio-emotional support. Vlahovic, Wang, Kraut, and Levine (2014) examined the match between support sought and obtained and found that a fit between the two was especially important for individuals seeking information.

Ren, Kraut, and Kiesler (2007) suggest that different modes of individual attachment have different implications for the general character of a community. For instance, relationship-based communities tend to have longer, meandering discussions that cover lots of topics, whereas identity-based communities tend to focus more narrowly on their joint task and the instrumental value of a platform. They also speculate that relationship-based communities are often not perceived as welcoming by newcomers even as they provide satisfaction to committed core members. Yet communities are not homogeneous, and people who derive different benefits from an online site may self-organize around an interlocking set of roles. Gleave et al. (2009) consider this level of organization as a community's "role ecology." For instance, *Answer* people in a Usenet technical support newsgroup require a relatively large population of *Asker* people to provide the questions that stimulate them to generate replies. Simultaneously, timely and useful replies from *Answer* people continue to attract *Asker* people to the site.

Butler's work (Butler, 2001; Butler, Bateman, Gray, & Diamant, 2014) illustrates how an exchange of resources among different subpopulations can produce dynamics in an online community, potentially leading toward a stable point where the pooled needs and resources of a population are in balance. However the technical design of an interface is a critical mediator in this systems-oriented picture (Butler et al., 2014). The *affordances* of a platform, referring to the way in which a user perceives that he or she can manipulate an artifact, influence how easily and what kinds of resources a user may produce or consume (Norman, 2002). For example, a micro-blogging platform does not readily afford the production of novellas or research articles. S*ignifiers*, on the other hand, carry information that users use to determine whether and how to engage with a platform (Norman, 2008). In social platforms, social signifiers such as the number of likes or comments a post receives can play an especially important role in peoples' decisions to engage with that content.

However, there remains a paucity of empirical research to examine the impact of design on the sustainability of OHCs. In one recent study, Nakikj and Mamykina (2018) used qualitative methods to highlight how a design change that was intended to facilitate information seeking behaviors interfered with the provision of socio-emotional support within the community, and this may have downstream consequences for the community's sustainability. However, their analysis focused on the perceptions of community members and did not follow the impact of the design change through to any observed member behaviors. With the work presented here we provide a complementary, systems-oriented perspective on a design change. Introne et al. (2016) hypothesized that core members in WebMD forums participated to maintain their relationships while providing newcomers with much of the information support they sought. We build on that work to consider whether this arrangement was likely to be sustainable, and how changes in the affordances and signifiers of the platform might have altered it. We organize our study around four research questions:

RQ1: *How did the distribution of posting traffic among core members and newcomers change across the design change*?

RQ2: *How did relationship-oriented activity change across the design change*?

RQ3: *How did changing social signifiers influence newcomers' decisions to post*?

RQ4: *How was the design change perceived by existing, and, in particular, core members*?

As necessary background we first provide a detailed analysis of WebMD's design before and after the change that took place on February 28, 2010.

DOI: 10.1002/asi

Research Setting

WebMD.com is a popular, health-related website. As of January 2016, Alexa[1] ranked the site as the 105th most popular in the United States and the second most popular health-related site behind NIH.gov (Introne et al., 2016). In addition to providing health-related news, resources aimed at assisting visitors with self-diagnosis information, and a directory of medical professionals, WebMD hosts an array of topic-specific message forums. These forums have a long history on the site and have been alternately referred to by WebMD as message boards, support groups, discussions, and exchanges. Throughout these name changes, the forums have continually been used to host conversations oriented around the exchange of social support.

In the WebMD forums, as we refer to them here, conversations are organized into threads. Although there are no technical constraints on what people post, most threads begin with a *request* for social support (implicit or explicit) followed by a series of *responses* that strive to fulfill the request. However, interlocutors also weave a variety of other kinds of interaction throughout these conversations, sharing details of their daily activities, humorous anecdotes, and other conversation that has little to do with the specific health condition that the forum is intended to focus on (Introne et al., 2016).

The basic information architecture (i.e., technical organization) of WebMD's forums has not changed since its initial inception. Each forum in WebMD is specific to a health condition (e.g., diabetes, ADHD) or topic (diet, raising children). In any given forum, conversations are organized into many discussion threads. A request post has both a short, user-assigned title, and a longer body; responses only have a body. Users can choose to respond to either the original request or to one of the previous responses.

WebMD elected to introduce a significant design change in the online forums at the end of February 2010 (contrast Figures 1 and 2). Although the precise intent behind WedMD's decision to evolve its branding, organization, and the look and feel of its forum sites is not known, several communications from WebMD around the time of the redesign provide a few clues. Roughly 2 weeks before the change, moderators posted a lengthy message in the forums describing the upcoming redesign. An excerpt of this message is included here:

> *Dear WebMD Members, we've been providing you with tidbits about the new features that are coming, and we're excited to finally announce the name of our new community service, The WebMD Health Exchange! The WebMD Health Exchange will offer you a highly interactive and enhanced community experience that will help bring you closer to more health and wellness experts as well as others like you. In addition, The WebMD Health Exchange uses structured Discussions, Tips & Resources to gain valuable insights from the community and surfaces the conversations members find "most valuable" by measuring activity within those conversations.*

[1]Alexa is a commercial tool used to rank various Internet sites. It can be considered a rough indicator of site popularity.

Before the design change, WebMD hosted 138 different condition specific forums, referred to on the site as message boards. Although the naming conventions varied (e.g., "Weight Loss Support: New Moms" "Stroke: Support Group," "Women's Health: Friends Talking,") 65 of the forum titles contained the word "support," 57 in the phrase "support group." Each of the message boards included a *thread index page* for each topic specific forum (shown in Figure 1), which offered the user an index of all discussion threads, sorted from the most recently active thread down (with the exception of threads that were pinned to the top by a moderator). A total of 25 threads were shown on a page, covering roughly 600 pixels of vertical space, easily fitting on a single screen in most displays in 2009. Each entry in the thread index provided information about the author who initiated the thread, the date of the last post, and the number of responses; the word "NEW" was present if there were no responses, or if the user was logged in and there were unread responses. Clicking the "+" to the left of a thread title displayed a list of all posts in the thread along with the post author and date of the post. If a user were logged in, the interface would display the word "NEW" in the rightmost column for any posts that were new to that user.

Posts belonging to each individual thread were displayed on a *thread detail page* (not shown here). Posts on the detail page were sorted first by reply order and then date order, with newer posts appearing later in the list. Up to 150 posts were shown per page, and the user could navigate to subsequent pages of posts using a page index. By clicking on the title of a thread in the thread index, users could navigate to the first post (the question) on the thread detail page. Alternatively, clicking on any of the reply indicators nested beneath the thread title on the thread index page (revealed by clicking the "+" button to the left of a thread title) would navigate directly to the reply on the thread detail page.

With the design change, the 138 message boards were replaced with 93 rebranded "health exchanges." The content from the older message boards was preserved, but some of the message boards were merged to become a single health exchange. For example, the message boards "Maintaining Weight Loss: Support Group," "Weight Loss Support: New Moms," "Weight Loss Surgery: Support Group," "Underweight: Support Group," "Diet Tool & Tips" were bundled together into the newly formed "WebMD Diet Exchange."

Fifty of the exchanges combined content from 92 of the previous message boards and were given a "featured" exchange designation, which indicated that they were moderated and cleaned of spam. Numerous medical experts (41) were also hired and assigned to answer member questions in some of the featured forums as part of the design change, although the allocation of experts to forums was uneven. These featured exchanges were uniformly named "WebMD <focus area> Exchange" and were given a prominent index page on the site. The remaining message boards were renamed as "<focus area> Member Exchange." These forums were member-managed and could be found by going to an index of all exchanges on the site. None of the names of the new exchanges contained the word "support." WebMD also

DOI: 10.1002/asi

FIG. 1. WebMD forum interface pre-February 2010: visitors are presented with a thread-based index that summarizes current topics and activity. Each entry under the heading "Discussion Title" is a thread, and the "+" icon to the left of each thread provides summary information about the users who have posted to the thread (screen capture from the internet archive). [Color figure can be viewed at wileyonlinelibrary.com]

introduced the ability for members to create new, member-managed exchanges if they wanted.

In addition to these organizational and branding changes, the redesign included significant revisions to the interface layout and functionality (refer to Figure 2). A detailed summary of these changes is presented in Table 1. One consequence of these changes was that it was harder for individuals to navigate long conversations and figure out who was present in a conversation before deciding to post.

Methods

Unlike Nakikj and Mamykina (2018), we did not have access to individual participants in the forum. To gather evidence supporting causal inferences about the impact of the design change, we examined four different aspects of the data, comparing different measures across the design change. Addressing RQ1, we visualized longitudinal posting patterns as distributed across different types of members in the community and in different communities. Addressing RQ2, we used an interrupted time series analysis to examine changes in the reply rate for people with stronger and weaker relationships. Addressing RQ3, we used regression modeling to draw inferences about the impact of different social signifiers on the rate of new posts made by first time posters. Addressing RQ4, we used content analysis to examine members' reactions to the design change and gauge their levels of frustration.

Before our analysis, we reviewed the terms of service on WebMD and carefully considered the ethical aspects of the endeavor. The terms of service explicitly allow data collection for noncommercial purposes, and text from user posts could be reproduced as long as the WebMD's copyright notice was included.[2] Users are informed by WebMD that there is no expectation of privacy in the forums, and the terms state the no personally identifying information may be posted. Personally, identifying information is removed when identified by site moderators.

Collection and analysis of these data was not considered to be Human Subjects Research by the authors' Institutional Review Boards, because it is public data, and neither

[2]Such uses are explicitly allowed in the WebMD terms of service: "WebMD authorizes you to view or download a single copy of the material on the WebMD Site solely for your personal, noncommercial use if you include the copyright notice located at the end of the material, for example: '©2013, WebMD, LLC. All rights reserved'"; http://www.webmd.com/about-webmd-policies/about-terms-and-conditions-of-use?ss=ftr - part3

DOI: 10.1002/asi

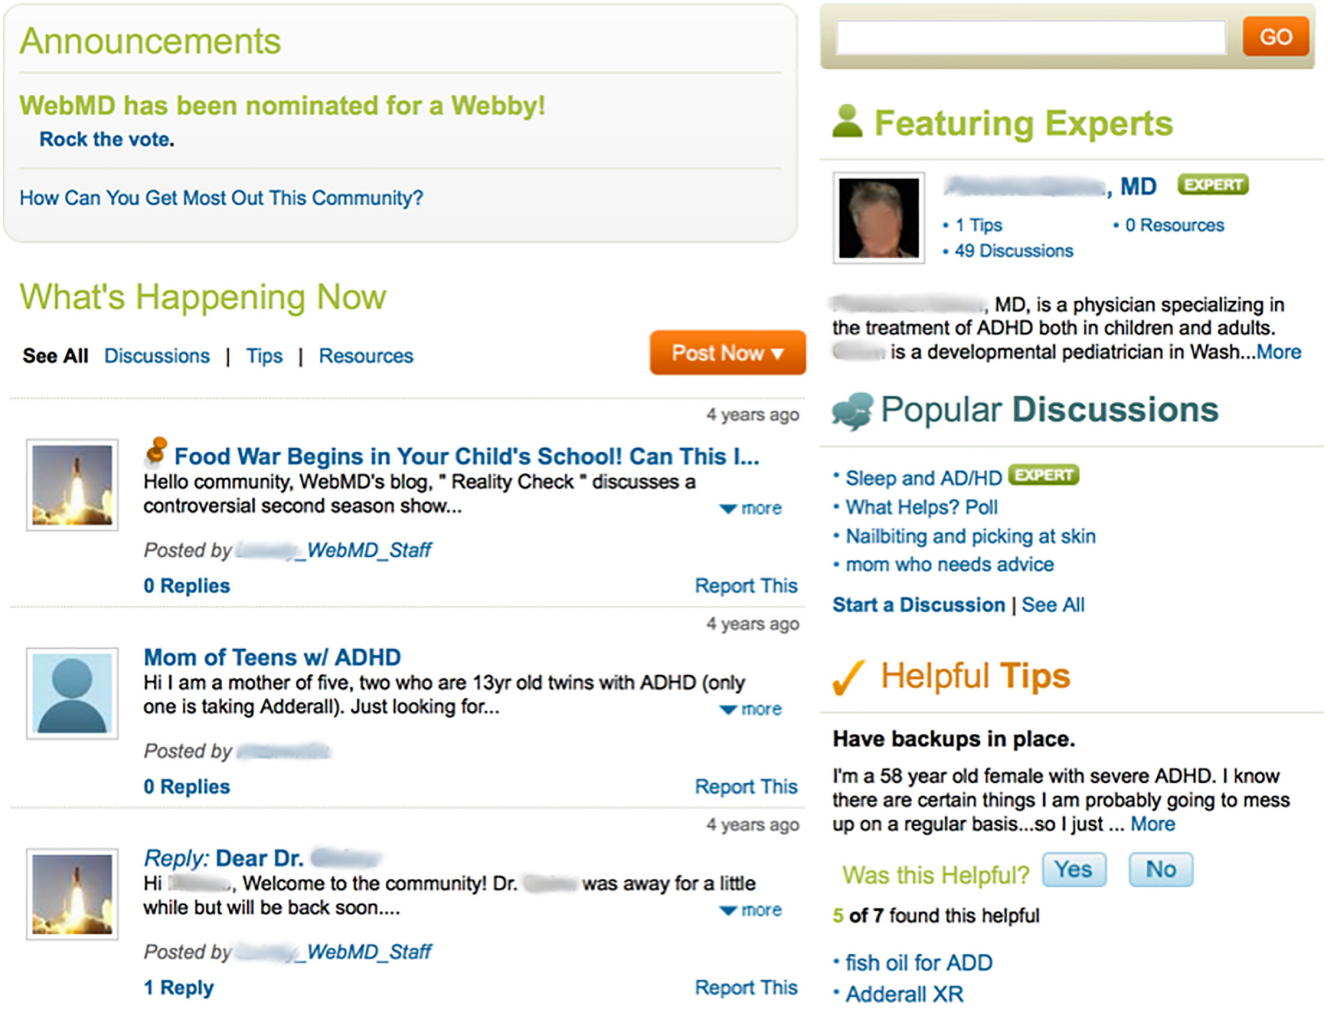

FIG. 2. WebMD forum interface post-February 2010: Visitors to the site are offered a digest with the most recent replies (screen capture from the internet archive). [Color figure can be viewed at wileyonlinelibrary.com]

introduces a manipulation nor involves any personally identifying information that would compromise forum members' privacy. However, we recognize that sensitive information might still be embedded within user posts. Thus, following conventions established in prior published work on WebMD (Huh, 2015; Huh et al., 2013), we only included quotes that we felt were critical to our argument and did not contain potentially sensitive information, and use pseudonyms in all posts.

In the following subsections, we first describe our data preparation methods, and then describe our analytical methods. To improve readability, we provide a brief overview of our methods here and refer readers to the online appendix (Appendix A: Detailed Methods) for additional detail.

Data Collection and Preparation

At the end of August 2014, the first author scraped all available content from the fifty-five featured forums on WebMD. This data consists of all messages that were available when the data was scraped, comprising a total of 1.1 M posts from roughly 275 K users spanning seven years (i.e., 2007 through 2014). After scraping the data, it was cleaned as described in prior work by Introne et al. (2016). Data before 2009 was incomplete, and so we further restricted our analysis to posting data from the period January 1, 2009 and ending August 1, 2014. We also discovered that five of the forums did not exist before the design change and another existed for only three months before it, so we omitted these six forums from the analysis. After restricting the data thusly, 1,079,575 posts generated by 214,252 posters remained for analysis.

For these 49 forums we inferred social network structures from reply patterns, creating a link, or "relationship," between any two users who had at least one dyadic interaction in a single thread. An examination of the social networks the those merged forums that were distinct before the design change led us to omit an additional forum (the Sex & Relationships Exchange) from subsequent statistical analyses. Finally, following Introne et al. (2016) we applied a core-periphery network analysis to classify users into core, peripheral, and extra-peripheral users.

DOI: 10.1002/asi

TABLE 1. Summary of design changes in WebMD; rightmost column indicates whether the change effected an affordance (A), a signifier (S), or both (A/S).

Feature	Before design change (pre-February 2010)	After design change (post-February 2010)	A/S
Context and navigation	Forums were referred to as message boards and were undifferentiated.	Forums referred to as "Exchanges" and were divided into "featured' and member exchanges.	S
	All forums accessible from a link on the home page.	Featured forums accessible from an "exchange" landing page. The full list of exchanges (including member exchanges) available on a separate index page.	A/S
	No consistent naming scheme; many forum titles contained the work "support"	Featured exchanges were named "WebMD <focus area > Exchange," other named "<focus area > Member Exchange." No forums contained the word support.	S
	Members cannot create new forums	Members can create new forums	A
Thread index page	Content is evaluated on a 5-star rating displayed after discussion title	No rating method	S
	New posts would be displayed as such for logged in users.	No notification of posts new to user	S
	25 threads shown by default	30 threads shown (not adjustable)	S
	Author of the thread is shown	Author of the last post is shown	S
	Clicking the "+" sign to left of thread title displayed list of all posts in thread and summary information about the users who have posted to the thread	Clicking the "more" link reveals brief snippets of text from the full post, but nothing about the posters who have contributed to the discussion.	A/S
Thread detail page	Profile pictures of post author are not displayed next to discussion title	Profile pictures of post author are displayed next to discussion title; badges for experts are shown.	S
	Displayed 150 posts per page	Displayed 15 posts per page	S
	User can navigate to specific post from the thread index	Users can only navigate to the first post in the thread from the index	A
	Profile pictures of post author are not displayed next to post	Profile pictures of post author are displayed next to post; badges for experts are shown.	S
	Messages are indented to show threading, and organized by replies	Messages are not indented, and are organized by date; however reply structure can be recovered by clicking on a link	A/S

Analysis

We organized our analytical methods around the preceding four research questions, which are repeated here for clarity.

RQ1: How did the distribution of posting traffic among core members and newcomers change?

We visualized longitudinal traffic rates to get a sense of general trends in the data. We then inspected the relative levels of two kinds of traffic over time: *first-time* posts, which include all posts made by newcomers, and *relationship* posts, which are replies made between dyads who have an existing link in our inferred social networks. Finally, we used a series of visualizations to understand changes in how activity was distributed across the population before and after the design-change. In these analyses, we distinguish between *top-level* posts, which initiate threads, and replies, which occur in response to other posts. We also examine these quantities in relation to two measures derived from our core-periphery analysis: *core size* refers to the number of individuals who are labeled as core, and *fitness* as a measure of how similar a network in a given forum is to and idealized core-periphery network (Borgatti & Everett, 2000). Fitness varies from (0–1), and higher degrees of fitness indicate that core members are more densely interconnected, and fewer peripheral members have relationships with one another.

RQ2: How did relationship-oriented activity change?

To understand how the design change impacted relationship maintenance activity, we contrasted changes in the weekly rate of posts between dyads who had strong relationships and those who did not across the design change. Here, we operationalize a strong relationship as one that has met a threshold (T) of a certain number of dyadic interactions. There is no theoretically motivated way to establish T, so we modeled our data across a range of values of T (from 1–10). Based on the appearance of an inflection point in posting trends around the design change, we performed our analysis using a mixed effects interrupted time-series analysis, allowing the intercept to vary randomly across the different forums.

RQ3: How did social signifiers influence newcomers' decisions to post?

To understand how the design change impacted newcomer traffic, we used regression analysis to examine how visible social signifiers influenced the rate at which new posts were made before and after the change. Here, we assume that a population of potential newcomers visits a forum at any given point time. The size of this population varies depending on the nature of the health condition and time (of day, of week, and possibly of year), but these visits are otherwise identically and independently distributed (IID). Whether or not a potential newcomer posts is

DOI: 10.1002/asi

based in part on the activity that is visible in the forum, and the design of the interface dictates how this information is presented.

To capture this in a regression model, we worked forward through our data to reconstruct the activity that would be visible in the thread index page in either design, and captured the information present in the visible signifiers (see Table 1) using a set of five variables:

- **NonResponseRate:** The proportion of visible threads without a response
- **ProportionStaff:** The proportion of threads with a visible WebMD staff member
- **ResponseTime:** Min/max/mean amount of time between the top-level request and first response, aggregated across threads.
- **LastActivity:** Min/max/mean amount of time elapsed since the last post in the thread, aggregated across threads.
- **ThreadLength:** Min/max/mean length of visible threads.

For the last three variables, we explored each of the aggregation methods (min/max/mean) separately in our regression models. All models also included several control variables: forum identify, an interaction term between the number of weeks from the beginning of the data and forum identity, weekday, and hour. Our outcome variable was the waiting time between each first-time post and the post immediately preceding it. Further discussion for the motivation for and derivation of each of these variables is provided in the appendix.

RQ4: How was the design change perceived by members?

To gain insight into members' reactions to the design change we used human coders to examine posts about the forum itself, which we refer to here as meta-posts. We did this in two rounds. First, we selected all posts in a two-month window centered on the design change (from Feb. 1, 2010 to Apr. 1, 2010), and then filtered those posts using MTurk workers to determine which posts were about the forum itself, and whether the disposition of the post was positive or negative. The latter was evaluated on a five-point likert scale, from −2 (i.e., "very negative") to 2 (i.e., "very positive"). This yielded a total of 1107 posts for further analysis.

In a second round, two experienced coders were asked to assess the level of frustration in each post, as well as classify the nature of the problems reported. To assess frustration, we developed a five-point coding scheme based on posters' indications they might leave the forum:

0. Post is a staff announcement (hence has no bearing on the likelihood of leaving).
1. User is unlikely to leave.
2. User is frustrated by the platform but unlikely to leave.
3. User is frustrated by the platform and likely to leave.
4. User is frustrated by the platform and was very likely to leave or has already left but is making a brief post.

We developed five categories to characterize the nature of the reported problems:

- **Error:** posts that mention errors or bugs in the WebMD platform
- **Layout**: posts referencing the WebMD platform layout and other related issues
- **Functionality:** posts describing the WebMD platform and system functions afforded to users
- **Administration:** posts that mention administrative actions on the part of WebMD employees or matters that required their action
- **Other:** posts that were about the forum but deemed irrelevant to determining why users might be inclined to stop using WebMD. Many such cases involved users telling one another how to use a specific feature of the platform

Results

Traffic Distributions

Figure 3 offers an overview of posting traffic in the 55 featured forums. Traffic patterns are normalized in Figure 3 to show detailed fluctuations in traffic, but forum activity levels vary dramatically, from approximately 170 K to 2300 total posts. The mean total volume of posts across the forums is 21,640, the median is 9,941, and the inter-quartile range is approximately 20 k. Despite significant heterogeneity across forums, there is an overall reduction in posting activity beginning at the point of the design change.

Figure 4 breaks longitudinal traffic rates into average weekly posting traffic for both first time and relationship posts. In many cases, there was an upswing in some kinds of activity following the design change. The increase was especially pronounced among relationship-oriented posts in the smaller forums. However, in almost all cases, traffic levels dropped off regularly in the years following. In the last eight-month period, two forums still had more weekly traffic among first-time posts and overall traffic than before the design change, but 19 of the smaller forums still exhibited more relationship posts.

Figure 5 visualizes the core-periphery structure underlying each forum in the first and last 14 months of the dataset, along with the ratio of top-level posts to all posts, labeled as the Request Ratio in Figure 5. Lower request ratios indicate that a forum is more responsive (because there are more responses per request) and that conversations are richer (because most conversations occur in within long threads). Before the design change, larger forums (as measured by total number of posts) tended to have larger cores ($R = .78$, $p < .0001$), often with high fitness. These forums also exhibited the lowest request ratios. By the final 14-month period, no forums have a core with more than 10 individuals. Several of the groups with smaller cores also increased in fitness, suggesting better connected cores and a reduction of conversations within the periphery.

Relationship-Oriented Activity

Our visual analysis suggested an inflection point in traffic trends around the design change, and we used an interrupted time series analysis to examine this observation

DOI: 10.1002/asi

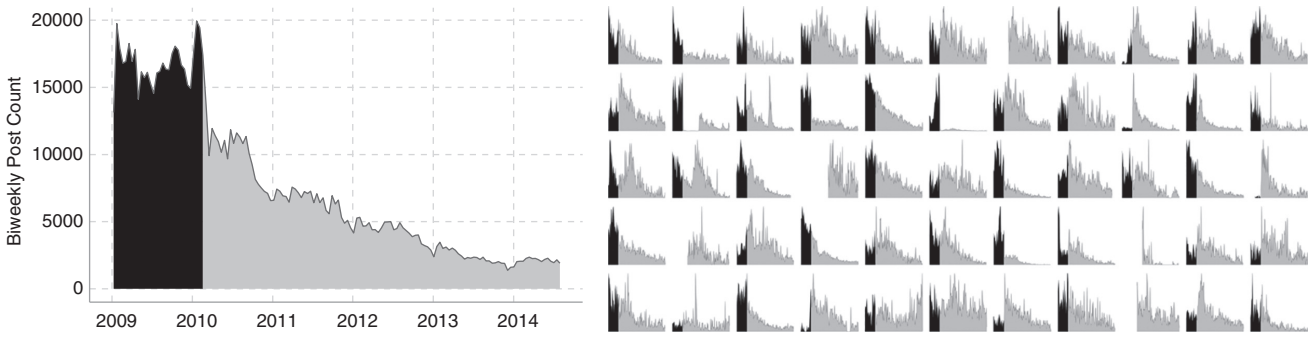

FIG. 3. Summary of traffic trends before and after the design change. Black shading is before the design change, gray is after. Traffic is aggregated biweekly. The image on the right is biweekly posting traffic aggregated over all forms. The image on the left provides an overview of the 55 forums; the *y*-axis in each case is normalized to the maximum period of traffic for that forum.

statistically. Given notable differences in core member activity and the relative number of replies, we compared trends in replies between dyads with strong relationships and those without, varying assumptions about how many interactions (T = [1–10]) constituted a strong relationship.

The results presented in Table 2 indicate that our models captured a good portion of the variance in the data, but the majority of this is owing to forum identity, especially in the case of relationship replies. Coefficients reflect the number of posts per week per forum, and there are roughly 60 weeks before the design change, and 213 following. Thus, when T = 1 the mean communication rate between dyads who had previously interacted starts at about 10.7 posts/week, grows to about 21 posts/week by the design change, and then (ignoring the nonsignificant factor) diminishes to about 8 posts/week by the end of the dataset.

There are many more nonrelationship-oriented posts than relationship posts, and in general these decay slightly before the design change, and then more quickly afterwards. Relationship posts, though less frequent, grow before the design change but then trend downward. The magnitude of the difference diminishes as the threshold T is increased, but such strong relationships are quite rare to begin with. For example, when T = 10, the number of relationship posts grows to about 3 per

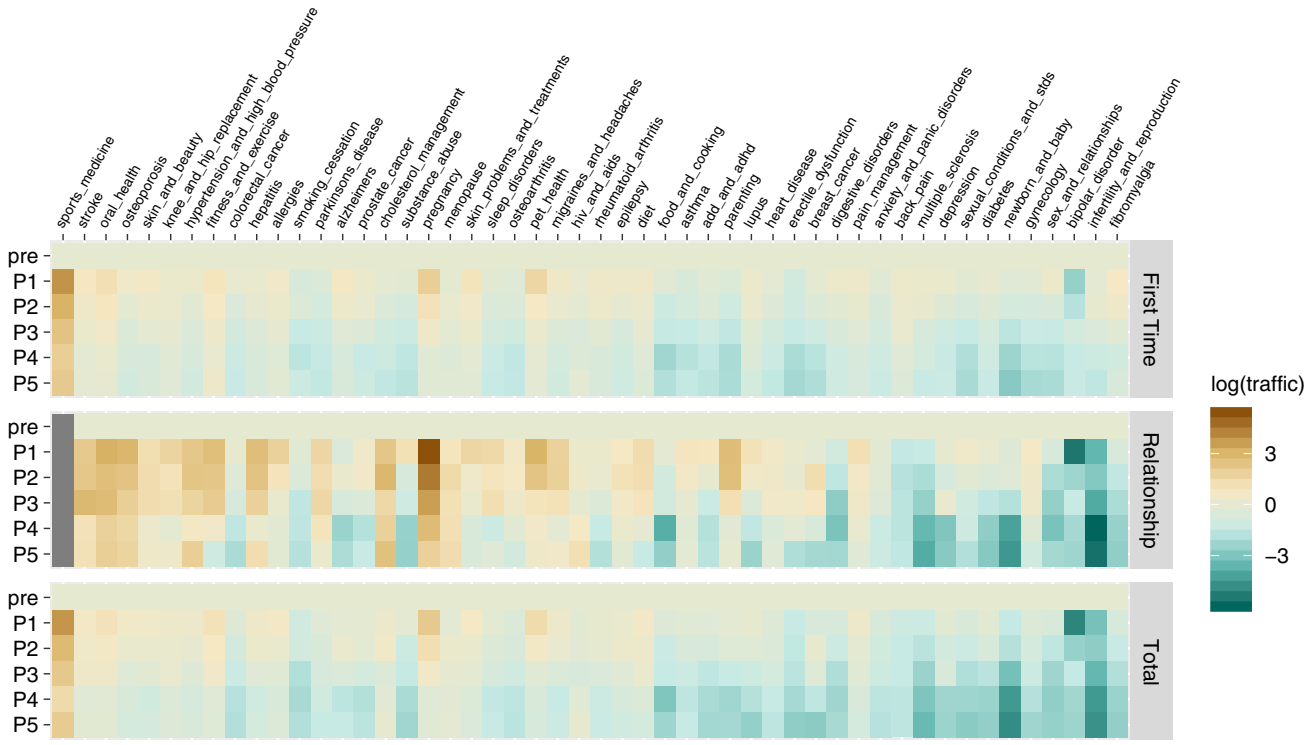

FIG. 4. Traffic rates for the 49 forums relative to that before the design change. Colors indicate average weekly traffic relative to the levels before the design change ("pre"). Periods (along the left axis) represent year-long intervals except for the first (14 months) and the last (8 months). Traffic is log scaled, such that any contraction is less than zero, and any growth is greater than zero. Forums are sorted from left to right in order of increasing activity in the 14-month period preceding the design change. Note that the "sports medicine" forum did not have any relationships as we have defined them here. [Color figure can be viewed at wileyonlinelibrary.com]

DOI: 10.1002/asi

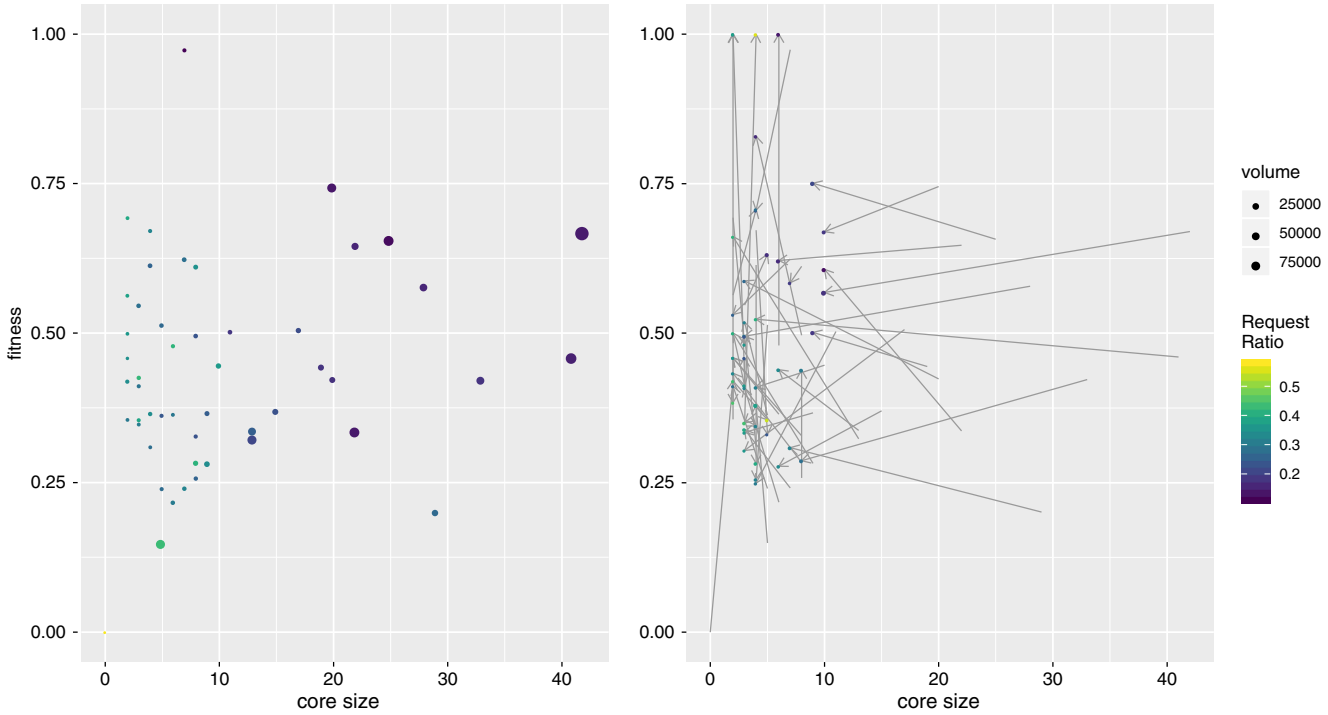

FIG. 5. Plot of change in core metrics between the 14-month period preceding the design change, and the final 14 months. Point sizes reflect the total number of posts in the period, and Request Ratio is the number of top-posts / the number of replies. [Color figure can be viewed at wileyonlinelibrary.com]

week, but quickly drops to more than half of that immediately after the design change, and then disappears completely by the end of the data.

Social Signifiers and the Decision to Post

Turning our attention to first-time posts, Table 3 summarizes our results for the best models describing the influence of social signifiers before and after the change. Both models explained roughly two thirds of the deviance in the data, but noncontrol factors accounted for about 86% of this in the prechange model, and only about 28% in the postchange model.

In both models, variables are expressed in terms of hours, so comparing their intercepts indicates that the average waiting time increased from about a half of an hour to just over an hour after the design change. In both models **MinLastActivity**, which is a visible indicator of how heavily trafficked a forum is, accounted for the most deviance because of noncontrol factors. Before the design change, this covariate explains more deviance than the other factors combined; for every additional hour since the last activity in a forum, average wait time increases by nearly 10 min. After the change, this factor explains a much smaller proportion of deviance, increasing wait times by about 5 min for every additional hour without a post. The influence of **AvgThreadLength** follows a similar pattern. An increase of one post in average thread-length reduces wait times by about 27 s before the change and 16 s after the change.

Before the change, **MinResponseTime**, which could be assessed in the interface by expanding the posts beneath a thread in the index, was also a significant factor. In this case, when the fastest response to a request was delayed by an hour, the wait time increased by about 1 min. This could have a large cumulative effect if response times are extended for multiple hours. After the design change minimum response time could only be assessed by viewing the thread detail page, and so it is not surprising that it was no longer predictive. However, the presence of staff in the thread index appears to dissuade new posters; each thread (of the 30 visible in the index) with a visible staff member increases wait time by about 1.3 min.

In summary, signifiers of high activity, including the recency of posts, the length of conversations, and (before

TABLE 2. Summary of modeling results for interrupted time series at different thresholds. "Rel" indicate coefficients for models that predict weekly active relationships at the given threshold value (T), and "N.Rel" coefficients predict nonrelationship activity at the same threshold. Both conditional and marginal pseudo-R^2 are reported.

	Rel.	N.Rel.	Rel.	N.Rel.
	T = 1		**T = 3**	
Intercept	10.68*	65.79***	2.71 n.s.	73.77***
DesignChange	−5.57 n.s.	−16.45 n.s	−4.73**	−17.39 n.s.
Time	.17***	−.16***	.11***	−.11***
TimeAfter	−.23***	.002 n.s	−.14***	−.10**
Pseudo R^2 Cond(Marg)	.76 (.05)	.80 (.14)	.71 (.03)	.80 (.12)
	T = 5		**T = 10**	
Intercept	1.42 n.s.	75.07***	.59 n.s.	75.89***
DesignChange	−3.25**	−18.87 n.s.	−1.8**	−20.32 n.s.
Time	.09***	−.08**	.05***	−.05 n.s.
TimeAfter	−.10***	−.13***	−.06***	−.17***
Pseudo R^2 Cond(Marg)	.69 (.02)	.80 (.12)	.66 (.02)	.80 (.12)

DOI: 10.1002/asi

TABLE 3. Comparing the statistical impact of the interface on first time posters before and after the design change. %Dev is the % of total deviance explained by the model. Both models used the same control factors: Hour, Weekday, Forum, Forum:WeekIndex, which account for the balance of deviance explained in both models. All effects are highly significant at $p < .0001$.

Waiting time before change			Waiting time after change		
Factor	Coeff	%Dev.	Factor	Coeff	%Dev.
Intercept	4.90×10^{-1}	—	*Intercept*	1.10	—
MinLastActivity	1.65×10^{-1}	84.70%	MinLastActivity	8.27×10^{-2}	27.79%
AvgThreadLength	-7.40×10^{-3}	0.91%	AvgThreadLength	-4.37×10^{-3}	.02%
MinResponseTime	1.71×10^{-2}	1.10%	ProportionStaff	6.68×10^{-1}	.01%
Adj. D ^2 = 66.57%			Adj. D ^2 = 63.73%		

the design change) the responsiveness of the forum are predictive of first-time posters' decisions to post. Before the change, these factors overwhelm other sources of deviance in the model, playing a larger role than temporal trends, seasonal variation, or even forum identity. These results suggest a strong coupling between visible relationship-oriented activity and the rate of newcomers. Following the change, this coupling is weakened. Activity matters less and responsiveness is no longer a factor. However, the influx of staff and medical experts that occurred in the design change may have in fact dampened newcomer traffic.

Reactions to the Design Change

We examined the content of posts to develop insights about how users experienced the design change and which aspects they did not like. One hundred twenty-one of the 1107 posts examined were found to be staff announcements, and we describe the remaining 986 here. Table 4 provides a numeric summary of these posts, and Figure 6 and Figure 7 additional details on the types of complaints made by different sorts of users. The core was responsible for most of the negative and frustrated posts following the design change, and these were centered on layout and functionality. This is not surprising; core members post more in general, and they were likely to have become accustomed to the interface before the design change.

Analyzing the content of these posts revealed the one common complaint about functionality voiced by core members was the difficulty keeping track of friends and conversational context in the new design. Several quotes from different core members help to make these frustrations clear. Note in the following, we include the coded frustration level:

TABLE 4. Numeric summary of complaints by user class; *Frust* indicates the assessed level of frustration of the poster, based on an expressed intention to leave.

	Before			After		
	Frust>2	valence (mean)	#	Frust>2	Valence (mean)	#
Core	2	.20	164	72	−.09	436
Periph.	5	.17	83	56	−.19	226
X-Periph.	1	.08	6	6	−.05	59
Totals	8	.18	265	134	−.12	721

"When we had the other board—we were able to click on the side and the list of posters would come up—you got to see who posted and when. Here you can't do that. If you want to read or know you have t click on it and then scroll down—then i get frustrated because it's the same posted answers—none added--but the only way to tell is to look at the whole post." **[Frustration: 2]**

"When I go to an exchange, and I scan the posts, I can only tell who the last person was who replied, not who originally posted it. So if I'm looking for an update from 'Bob', I have to open every post to see the original poster. Or if I tend to avoid a particular person (which I haven't, but you never know), I don't have to waste my time. I guess what I'm getting at is before we could see who posted originally, then hit the plus sign to see who responded. Now we can only see who last responded. It would be awesome if we could go back to seeing, per post, who the original poster is, on the front page." **[Frustration: 2]**

"This new way of posting makes it very hard for me to connect as much. I don't know what I've read, and what I haven't read, and just feel the posts are spread out too much. I like to see where I left off. In this new way, I have to read through the posts again. It's taking me too much time. I don't like being at the computer this much. This is just my opinion. So, if you don't see me as much, that's why. I love this site, but it's more work for me." **[Frustration: 3]**

These complaints offered very specific commentary that, as we discuss below, we believe is central for understanding how the design may have impaired relationship maintenance among forums with established cores.

Discussion

Before the design change, WebMD hosted a robust set of communities. There was substantial heterogeneity among those examined here, but a significant minority had robust networks of core members. As discussed by Introne et al. (2016), core members were responsible for a large proportion of the informational support delivered to newcomers and generated large volumes of more conversational traffic. Our results illustrate that such conversational interaction was in fact growing leading up to the design change. To newcomers, these forums would have appeared to have

DOI: 10.1002/asi

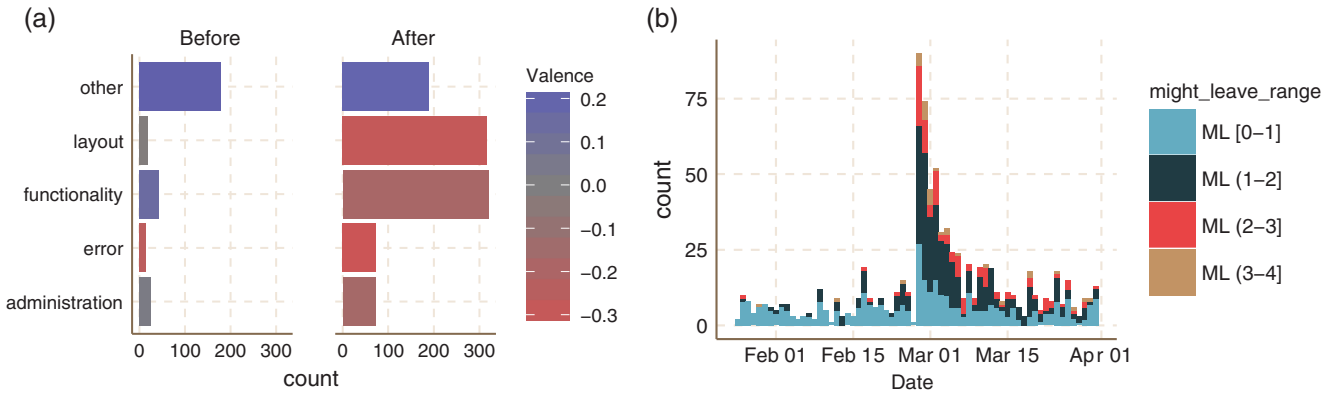

FIG. 6. (a,b): Visual summary of complaints before and after the design change. Graphs are based on 986 posts with a meta score > = 2 and exclude staff announcements in the two-month period centered on the design change. The graph on the left (a) indicates the number and type of complaints before and after the design change, colored by mean valence (note that there were multiple complaints in some posts, and 1250 complaints are represented). The graph on the right (b) is the number of posts aggregated daily and colored by number of posts in each indicated range of "might leave" scores. [Color figure can be viewed at wileyonlinelibrary.com]

many long conversations, highly responsive users, and regular activity. The appearance of this activity increased the likelihood that newcomers would post.

The design change altered the affordances and signifiers of the forum in various ways. We cannot be certain of WebMD's intent or design practices, but their archived messaging appears to have emphasized the quality of medical information over community building. The reactions of core members indicate that the new design reduced support for the community maintenance activities they felt were important. In particular, the lack of facilities for keeping track of people in long conversations seemed particularly upsetting for some. These results are reminiscent to those of Nakikj and Mamykina (2018), who report similar member reactions to a design-change that was oriented to improve information-seeking behavior.

Immediately following the design change, many forums experienced a brief bump in both first-time posters and relationship-oriented traffic, but both kinds of traffic subsequently decayed. Cores became smaller and posted less, and conversations grew shorter and less active. The design change appears to have somewhat obscured these signals of decay for newcomers, but the increased presence of staff had a negative impact on their decisions to post.

Adopting a systems perspective, we propose that the forums on WebMD before the design change might be

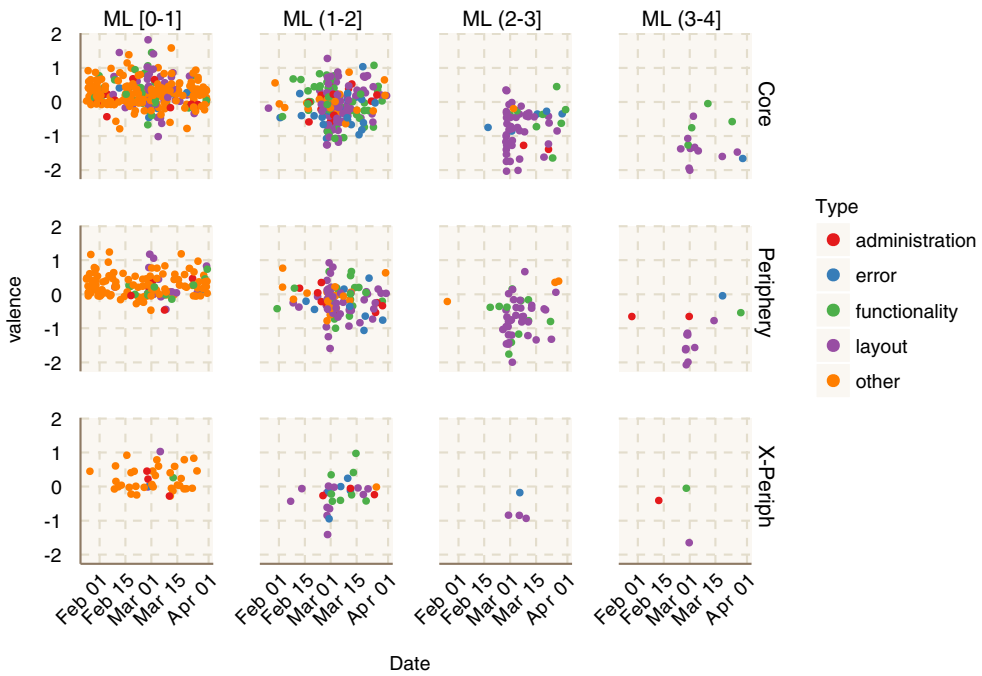

FIG. 7. Plot of complaint type by valence (*y*-axis), might leave range (columns), and poster class (rows) over time (*x*-axis). Covers the same data as in Figure 4; 1250 complaints are plotted. [Color figure can be viewed at wileyonlinelibrary.com]

DOI: 10.1002/asi

described as a coupled set of feedback loops that enabled it to maintain its levels of membership (see Figure 8). The first feedback process (on the right-hand side of Figure 8) reflects newcomer traffic, which is promoted via by the perception among newcomers that the forum is a good fit for their needs (Butler et al., 2014). This leads them to post requests, which in turn motivate responses by existing members. The second feedback loop (seen on the left-hand side of Figure 8) relates to core members of the community. In the context of WebMD, some site visitors create social bonds with others, becoming part of interconnected cores. These core members devote a significant amount of their time answering requests from incoming users, but also engage in community support exchanges through relationship posts (Bambina, 2007; White & Dorman, 2001). Indeed, previous studies suggest that the affordances and signifiers that manifest in relationship seeking behavior correlate more strongly with the use of a discussion board than information seeking (Chung, 2014). The two feedback loops are coupled, because core activity increases newcomers' perceptions that a forum is a good fit (e.g., they are likely to get a timely response), and newcomer questions create an opportunity for core members to respond and interact.

Following the design change, both the levels and distribution of posting activity changed, with notable decay in the size and activity of the core members. Across the design change, we propose that the system shifted from that depicted in Figure 8 to that in Figure 9. In the new system, the core feedback loop (left in Figure 9) is absent and is replaced by hired staff. Our results indicate that a visible staff presence in the thread index dissuades new posters, and it is therefore possible that while the system is in this state, traffic levels in the forums will continue to decline.

A critical element absent from our analysis is whether cores dwindled *because* of the design change, but the evidence we have gathered strongly suggests that this is the case. The design change eliminated the site affordances and signifiers that were important for the creation and maintenance of the rich network of relationships underlying the core. Moreover, the cores were not especially large to begin with—for example, the diabetes core, which was one of the larger before the design change, only had 25 members (Introne et al., 2016). Our best hypothesis is that some core individuals were motivated to leave by the design change itself, and others simply moved on as their circumstances changed. In either case, we further suspect that they were not replaced because the lack of support for community building. The other couplings that we have provided evidence for (for example between site activity generated by the core and newcomers' decisions to post, or between the visible presence of staff and newcomer's posts) are sufficient to predict the decay of the system in Figure 8 in the absence of a strong core.

Regardless of the causal impact of the design change, our analysis illustrates how a sociotechnical system can shift from one configuration that may be sustainable to one that may not be. Not all forums had arrived at the same point of sustainability before the design change—many of the smaller forums had very little relationship-oriented traffic. However, among those that had developed large cores, minor fluctuations were unlikely to disrupt the overall system. After the design change, without the presence of large cores, the system depends more heavily on paid staff to respond to newcomers, and in this sense is no longer sustainable.

Toward a Complex Adaptive Systems Framing of Online Communities

We believe that a useful framing for our analysis is to understand an online community as a system—in particular a complex adaptive system (CAS; Eysenbach et al., 2002; Holland, 2006). CAS theory characterizes systems to be composed of multiple, interacting, adaptive agents that generate large-scale emergent patterns in the absence of tightly centralized control structures. A CAS framing of a sociotechnical system—such as an OHC—focuses on the coupled interactions of agents in the system. Such couplings can be tight or loose, but they constrain the way an extended sequence of such interactions evolves. At scale, many sequences of interaction can be understood as feedback processes that are embedded in the system. Feedback processes can be one of two types: *reinforcing processes*, which accelerate or amplify the process itself, and *balancing processes*, which seek a goal-state (Senge, 2006, p. 79). For example, a chemical reaction that generates heat,

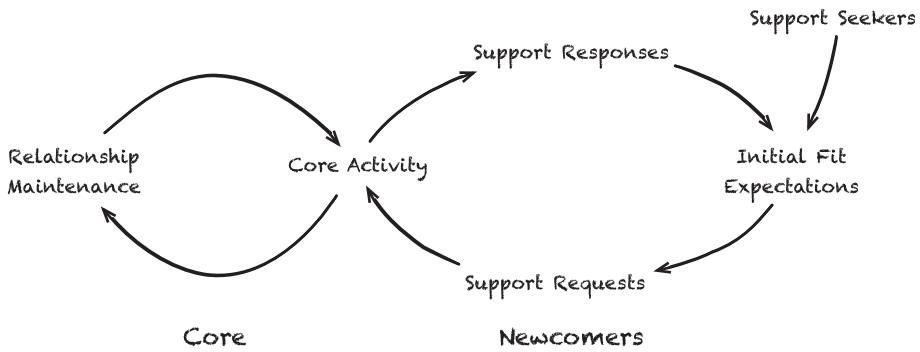

FIG. 8. The proposed set of coupled feedback loops operating before the design change.

DOI: 10.1002/asi

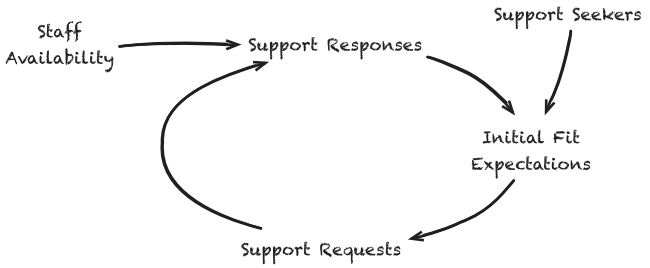

FIG. 9. The proposed system operating on WebMD following the design change.

which in turn accelerates the reaction itself, is a reinforcing process. If, however, temperatures that are too high slow the chemical reaction, the reaction becomes a balancing process. Finally, there may be limiting conditions that further constrain the system, such as when the reagents in a chemical reaction have been completely consumed.

Feedback processes within a complex adaptive system are revealed by the patterned sequences of interactions of agents over time (Kauffman, 1993). For example, the status of members in a team (a global variable in Arrow, McGrath, and Berdahl's [2000] treatment) emerges as a result of the history of their interactions. In some cases, the relative status of team members can begin to oscillate in regular patterns as sequences of interactions replay themselves (Arrow et al., 2000). Such patterns are called *attractors*, because they tend to attract or converge similar sequences of interactions. The construct of "attractor" in CAS operationalizes the tendency of a social system to evolve toward regular patterns of interaction across prominent social structures. In other words, participants in a stable online health forum move toward prominent, visible social groupings and both inter- and intragroup interactions exhibit regularity when examined at an aggregate level.

Any given attractor is associated with a *basin of attraction* that effectively bounds the amount of perturbation a system can withstand and still return to that attractor. As an analytical frame, attractors (and their basins) serve to reduce the complexity of dynamic systems because they drift toward attractors, and the number of attractors can be dramatically smaller than the number of possible interaction sequences (Kauffman, 1993, p. 193). The complete set of attractors defined by a dynamic system is its *attractor landscape*. Some exogenous changes to a system—related to a design change, or a shift in demographics—may be large enough to move a system into the basin of a different attractor, or may even alter the underlying attractor landscape itself. Because different attractors are characterized by different likely interactions among individuals, the benefit any individual is likely to get from a sociotechnical system depends on the attractor it is operating within. Importantly, not all attractors provide the same level of benefit to their members, and not all communities find stable attractors.

From this vantage point, some of the forums on WebMD moved from a stable attractor before the design to another that may or may not be stable afterwards. We believe that by manipulating the site's signifiers and affordances, WebMD's design change altered the attractor landscape in a manner such that the initial stable attractor no longer existed. Whether or not the second attractor is stable remains to be seen, but if so it is only through the continued infusion of resources by WebMD.

Conclusion

This research makes a major contribution to our understanding of the relationship between the design and stability of sociotechnical systems by showing how WebMD, a popular place for individuals to exchange social support about a variety of health conditions, maintained—or failed to maintain—its constituent set of rich and engaged communities in the face of a design change in 2010. Understanding the impact of this design change provides a key opportunity to examine the complex interrelationship between community participation, structure, and a site's design affordances.

In an era where increasing amounts of data and a renewed interest in AI are beginning to take center stage in the race to advance medical care, the simple and powerful activity of people helping people is often neglected. Yet, it is increasingly relevant. As individuals age and become less mobile, direct social support becomes less available, but the population of technologically savvy individuals will continue to grow. Online support platforms may be able to cheaply and effectively address the support needs of large numbers of people in ways improved algorithms and existing healthcare infrastructures cannot. They have the potential to make an appreciable difference in the lives of the people who use them. Yet, few studies have focused on how the design of these sites impacts the exchange of social support or the health of the communities themselves.

Against this backdrop, we see an urgent need to build advanced, sustainable systems that effectively balance the needs for high quality information, emotional support, and companionship. As we have seen, these systems are complex and dynamic as well as being prone to well-intentioned design innovations capable of unintended destruction. We believe the case of WebMD helps to identify a gap in existing design methodologies. Traditional design processes (e.g., user-centered design, value-centered design) focus on the needs of individuals, but often fail to consider the complex adaptive sociotechnical system many such interacting individuals create. In such venues, systemic factors will influence not only how well individuals' needs are met, but also will determine the ultimate stability of the system in the long term.

Acknowledging many remaining gaps in our knowledge, the insights articulated in this paper lead us to propose the importance of future research on the use of modeling as a strategic tool for developing sustainable sociotechnical platforms in the future. In the case of online support, models would allow designers to estimate how large a core must be to sustain a periphery of a given size or buffer the inevitable loss of core members. Perhaps

DOI: 10.1002/asi

more importantly, a model enables us to explore other configurations that might also be stable and have desirable properties. Finally, the development of such a representation brings additional research questions into focus that simply cannot be answered satisfactorily given available data, which would help initiate a virtuous cycle of modeling, research, and design to arrive at more innovative, sustainable online support systems. The work presented here helps point the way to advancing the science and art of designing sustainable systems for social support.

References

Ackerman, M.S. (2000). Intellectual challenge of CSCW: The gap between social requirements and technical feasibility. Human Computer Interaction, 15(2), 179–203.

Adams, S.A. (2010). Revisiting the online health information reliability debate in the wake of "web 2.0": An inter-disciplinary literature and website review. International Journal of Medical Informatics, 79(6), 391–400. https://doi.org/10.1016/j.ijmedinf.2010.01.006

Ahuja, M.K. & Galvin, J.E. (2003). Socialization in virtual groups. Journal of Management, 29(2), 161–185.

Arguello, J., Butler, B.S., Joyce, E., Kraut, R., Ling, K.S., Rosé, C., & Wang, X. (2006). Talk to me: Foundations for successful individual-group interactions in online communities. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 959–968). New York, NY: ACM.

Arrow, H., McGrath, J.E., & Berdahl, J.L. (2000). Small groups as complex systems: Formation, coordination, development, and adaptation. Thousand Oaks, CA: SAGE.

Bambina, A. (2007). Online social support: The interplay of social networks and computer-mediated communication. Amherst, NY: Cambria Press.

Berland, G.K., Elliott, M.N., Morales, L.S., Algazy, J.I., Kravitz, R.L., Broder, M.S., … McGlynn, E.A. (2001). Health information on the internet: Accessibility, quality, and readability in English and Spanish. JAMA, 285(20), 2612–2621. https://doi.org/10.1001/jama.285.20.2612

Borgatti, S.P. & Everett, M.G. (2000). Models of core/periphery structures. Social Networks, 21(4), 375–395. https://doi.org/10.1016/S0378-8733(99)00019-2

Burke, M., Joyce, E., Kim, T., Anand, V., & Kraut, R. (2007). Introductions and requests: Rhetorical strategies that elicit response in online communities. In Communities and technologies 2007 (pp. 21–39). London: Springer.

Burke, M., Kraut, R., & Joyce, E. (2010). Membership claims and requests: Conversation-level newcomer socialization strategies in online groups. Small Group Research, 41(1), 4–40. https://doi.org/10.1177/1046496409351936

Butler, B.S. (2001). Membership size, communication activity, and sustainability: A resource-based model of online social structures. Information Systems Research, 12(4), 346–362. https://doi.org/10.1287/isre.12.4.346.9703

Butler, B.S., Bateman, P.J., Gray, P.H., & Diamant, E.I. (2014). An attraction-selection-attrition theory of online community size and resilience. MIS Quarterly, 38(3), 699–728.

Carron-Arthur, B., Ali, K., Cunningham, J.A., & Griffiths, K.M. (2015). From help-seekers to influential users: A systematic review of participation styles in online health communities. Journal of Medical Internet Research, 17(12), e271. https://doi.org/10.2196/jmir.4705

Cheung, C.M.K., & Lee, M.K.O. (2009). Understanding the sustainability of a virtual community: Model development and empirical test. Journal of Information Science, 35(3), 279–298. https://doi.org/10.1177/0165551508099088

Choi, B., Alexander, K., Kraut, R.E., & Levine, J.M. (2010). Socialization tactics in wikipedia and their effects. In Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work (pp. 107–116). New York, NY: ACM.

Chung, J.E. (2014). Social networking in online support groups for health: How online social networking benefits patients. Journal of Health Communication, 19(6), 639–659. https://doi.org/10.1080/10810730.2012.757396

Civan, A., McDonald, D.W., Unruh, K.T., & Pratt, W. (2009). Locating patient expertise in everyday life. In Proceedings of the ACM 2009 International Conference on Supporting Group Work (pp. 291–300). New York, NY: ACM. Retrieved from http://dl.acm.org/citation.cfm?id=1531718

Costello, K.L., Martin, J.D., & Edwards Brinegar, A. (2017). Online disclosure of illicit information: Information behaviors in two drug forums. Journal of the Association for Information Science and Technology, 68 (10), 2439–2448. https://doi.org/10.1002/asi.23880

Coulson, N.S. (2005). Receiving social support online: An analysis of a computer-mediated support group for individuals living with irritable bowel syndrome. Cyberpsychology & Behavior, 8(6), 580–584. https://doi.org/10.1089/cpb.2005.8.580

Crowston, K., Wei, K., Li, Q., & Howison, J. (2006). Core and periphery in free/libre and open source software team communications. In Proceedings of the 39th Annual Hawaii International Conference on System Sciences, 2006. HICSS'06 (Vol. 6, pp. 118a–118a). Retrieved from http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=1579526

Csermely, P., London, A., Wu, L.-Y., & Uzzi, B. (2013). Structure and dynamics of core/periphery networks. Journal of Complex Networks, 1 (2), 93–123. https://doi.org/10.1093/comnet/cnt016

Cucuringu, M., Rombach, M. P., Lee, S. H., & Porter, M. A. (2014). Detection of core-periphery structure in networks using spectral methods and geodesic paths. ArXiv:1410.6572 [Cond-Mat, Physics:Physics]. Retrieved from http://arxiv.org/abs/1410.6572

Erfani, S.S., Abedin, B., & Blount, Y. (2017). The effect of social network site use on the psychological well-being of cancer patients. Journal of the Association for Information Science and Technology, 68(5), 1308–1322. https://doi.org/10.1002/asi.23702

Eysenbach, G., Powell, J., Englesakis, M., Rizo, C., & Stern, A. (2004). Health related virtual communities and electronic support groups: Systematic review of the effects of online peer to peer interactions. BMJ, 328(7449), 1166. https://doi.org/10.1136/bmj.328.7449.1166

Eysenbach, G., Powell, J., Kuss, O., & Sa, E.-R. (2002). Empirical studies assessing the quality of health information for consumers on the world wide web: A systematic review. JAMA, 287(20), 2691–2700. https://doi.org/10.1001/jama.287.20.2691

Fisher, D., Smith, M., & Welser, H. T. (2006). You are who you talk to: Detecting roles in usenet newsgroups. In Proceedings of the 39th Annual Hawaii International Conference on System Sciences, 2006. HICSS '06 (Vol. 3, pp. 59b–59b). https://doi.org/10.1109/HICSS.2006.536

Fox, S., & Purcell, K. (2010). Chronic disease and the internet. Washington, DC: Pew Internet & American Life Project Retrieved from http://www.pewinternet.org/files/old-media/Files/Reports/2010/PIP_Chronic_Disease_with_topline.pdf

Garcia, D., Mavrodiev, P., & Schweitzer, F. (2013). Social resilience in online communities: The autopsy of friendster. In Proceedings of the First ACM Conference on Online Social Networks (pp. 39–50). New York, NY: ACM. Retrieved from http://dl.acm.org/citation.cfm?id=2512946

Gill, P.S. & Whisnant, B. (2012). A qualitative assessment of an online support community for ovarian cancer patients. Patient Related Outcome Measures, 3, 51. https://doi.org/10.2147/PROM.S36034

Gleave, E., Welser, H. T., Lento, T. M., & Smith, M. A. (2009). A conceptual and operational definition of "social role" in online community. In 42nd Hawaii International Conference on System Sciences, 2009. HICSS '09 (pp. 1–11). https://doi.org/10.1109/HICSS.2009.6

Gorbatai, A. D. & Piskorski, M. (2012). Social structure of contributions to Wikipedia (working paper). Harvard Business School. Retrieved from http://www.wjh.harvard.edu/~hos/papers/AndreeaGorbattai/AndreeaGorbattai.pdf

Halfaker, A., Geiger, R.S., & Terveen, L.G. (2014). Snuggle: Designing for efficient socialization and ideological critique. In Proceedings

of the SIGCHI Conference on Human Factors in Computing Systems (pp. 311–320). New York, NY: ACM.
Halfaker, A., Keyes, O., & Taraborelli, D. (2013). Making peripheral participation legitimate: Reader engagement experiments in Wikipedia. In Proceedings of the 2013 Conference on Computer Supported Cooperative Work (pp. 849–860). New York, NY: ACM.
Hartzler, A. & Pratt, W. (2011). Managing the personal side of health: How patient expertise differs from the expertise of clinicians. Journal of Medical Internet Research, 13(3), e62.
Holland, J.H. (2006). Studying complex adaptive systems. Journal of Systems Science and Complexity, 19(1), 1–8. https://doi.org/10.1007/s11424-006-0001-z
Huh, J. (2015). Clinical questions in online health communities: The case of "see your doctor" threads. In Proceeding of the 2015 Conference on Computer Supported Cooperative Work. Vancouver, Canada: ACM Press.
Huh, J., McDonald, D.W., Hartzler, A., & Pratt, W. (2013). Patient moderator interaction in online health communities. In AMIA Annual Symposium Proceedings (Vol. 2013, pp. 627–636). Bethesda, MD: American Medical Informatics Association.
Introne, J., Semaan, B., & Goggins, S. (2016). A sociotechnical mechanism for online support provision. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (pp. 3559–3571). New York, NY: ACM. https://doi.org/10.1145/2858036.2858582
Kabadayi, S. & Price, K. (2014). Consumer – Brand engagement on Facebook: Liking and commenting behaviors. Journal of Research in Interactive Marketing, 8(3), 203–223. https://doi.org/10.1108/JRIM-12-2013-0081
Kaplan, B.H., Cassel, J.C., & Gore, S. (1977). Social support and health. Medical Care, 15(5), 47–58.
Kauffman, S.A. (1993). The origins of order: Self-organization and selection in evolution. New York: Oxford University Press.
Kraut, R.E., Burke, M., Riedl, J., & Resnick, P. (2012). The challenges of dealing with newcomers. In R. E. Kraut & P. Resnick (Eds.), Building successful online communities (p. 179). Cambridge, MA: MIT Press. https://doi.org/10.7551/mitpress/8472.003.0006
Lampe, C., Wash, R., Velasquez, A., & Ozkaya, E. (2010). Motivations to participate in online communities. New York, NY: ACM Press. https://doi.org/10.1145/1753326.1753616
Maloney-Krichmar, D., & Preece, J. (2005). A multilevel analysis of sociability, usability, and community dynamics in an online health community. ACM Transactions on Computer-Human Interaction (TOCHI), 12(2), 201–232. https://doi.org/10.1145/1067860.1067864
Mockus, A., Fielding, R.T., & Herbsleb, J.D. (2002). Two case studies of open source software development: Apache and Mozilla. ACM Transactions on Software Engineering and Methodology, 11(3), 309–346. https://doi.org/10.1145/567793.567795
Nakikj, D., & Mamykina, L. (2018). Lost in migration: Information management and community building in an online health community. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (pp. 146:1–146:14). New York, NY: ACM. https://doi.org/10.1145/3173574.3173720
Nambisan, P., Gustafson, D.H., Hawkins, R., & Pingree, S. (2016). Social support and responsiveness in online patient communities: Impact on service quality perceptions. Health Expectations, 19(1), 87–97. https://doi.org/10.1111/hex.12332
Norman, D.A. (2002). The design of everyday things. New York, NY: Basic Books.
Norman, D.A. (2008). THE WAY I SEE IT: Signifiers, not affordances. Interactions, 15(6), 18–19. https://doi.org/10.1145/1409040.1409044
Qureshi, I. & Fang, Y. (2011). Socialization in open source software projects: A growth mixture modeling approach. Organizational Research Methods, 14(1), 208–238. https://doi.org/10.1177/1094428110375002
Ren, Y., Kraut, R., & Kiesler, S. (2007). Applying common identity and bond theory to design of online communities. Organization Studies, 28 (3), 377–408. https://doi.org/10.1177/0170840607076007
Rodgers, S. & Chen, Q. (2005). Internet community group participation: Psychosocial benefits for women with breast cancer. Journal of Computer-Mediated Communication, 10(4). https://doi.org/10.1111/j.1083-6101.2005.tb00268.x
Senge, P.M. (2006). The fifth discipline: The art and practice of the learning organization. New York, NY: Doubleday/Currency.
Shim, M., Cappella, J.N., & Han, J.Y. (2011). How does insightful and emotional disclosure bring potential health benefits?: Study based on online support groups for women with breast cancer. The Journal of Communication, 61(3), 432–464. https://doi.org/10.1111/j.1460-2466.2011.01555.x
Stvilia, B., Mon, L., & Yi, Y.J. (2009). A model for online consumer health information quality. Journal of the American Society for Information Science and Technology, 60(9), 1781–1791. https://doi.org/10.1002/asi.21115
Taylor, S.E. (2011). Social support: A review. In H.S. Friedman (Ed.), The Oxford handbook of health psychology (pp. 189–214). Oxford, UK: Oxford University Press.
Vlahovic, T.A., Wang, Y.-C., Kraut, R.E., & Levine, J.M. (2014). Support matching and satisfaction in an online breast cancer support community. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (pp. 1625–1634). New York, NY: ACM. https://doi.org/10.1145/2556288.2557108
Krogh, G. von, Spaeth, S., & Lakhani, K. R. (2003). Community, joining, and specialization in open source software innovation: A case study. Amsterdam, Netherlands: Elsevier.
Wang, Y.-C., Kraut, R., & Levine, J.M. (2012). To stay or leave?: The relationship of emotional and informational support to commitment in online health support groups. In Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work (pp. 833–842). New York, NY: ACM. https://doi.org/10.1145/2145204.2145329
Welser, H.T., Gleave, E., Fisher, D., & Smith, M. (2007). Visualizing the signatures of social roles in online discussion groups. Journal of Social Structure, 8.
White, M. & Dorman, S.M. (2001). Receiving social support online: Implications for health education. Health Education Research, 16(6), 693–707. https://doi.org/10.1093/her/16.6.693
Yukawa, J. (2006). Co-reflection in online learning: Collaborative critical thinking as narrative. International Journal of Computer-Supported Collaborative Learning, 1(2), 203–228. https://doi.org/10.1007/s11412-006-8994-9