The Diffusion of a Task Recommendation System to Facilitate Contributions to an Online Community

Abstract

This paper studies the diffusion of WorkFinder, an intelligent task recommendation system, in Wikipedia. We investigate both factors that predict who adopts WorkFinder and its impact on adopters' future contributions to this online community. Analyzing logged data about participants' activities in Wikipedia, we found that both individual characteristics and social ties influence adoption. Specially, we found that highly involved contributors were more likely to adopt WorkFinder; and interpersonal exposure to innovation, cohesion, and tie homophily all substantially increased the likelihood of adoption. However, connections to admin contributors, i.e. those more prominent contributors in Wikipedia community, did not influence adoption. Finally, although the WorkFinder innovation saw limited distribution, adopters made significantly more contributions to Wikipedia after adoption than their non-adopter counterparts in the comparison group. (124 words).

Keywords: diffusion of innovation, network relationships, on-line community

The Diffusion of a Task Recommendation System to Facilitate Contributions to an Online Community

All communities, online and off, seek to motivate members to participate and continue contributing to the betterment of the group (Olson, 19xx; Kanter, 1972; Hechter 1988). Whether posting messages, welcoming newcomers, building information databases, or helping to administrate the group's policy, motivating people to contribute to online communities is essential for their sustained growth[1]. While many new communities are formed every day, many existing ones face the problem of under-contribution and/or no activities over extended period of time (Cummings, Butler, & Kraut, 2002; Ling et al., 2005). Even in active communities, the levels of contribution among participants can be extremely uneven. For instance, in open-source development communities, Lakhani and von Hippel (2003) found that 4% of members contributed 50% of the answers on a user-to-user help site. Mockus, Fielding and Andersen (2002) found that 4% of developers contributed 88% of new code and 66% code fixes. In our cross sectional data we found that during a randomly selected 28-day data collection window for this study, 10% of the 6,570 randomly selected participants did not edit any Wikipedia content at all. In contrast, the most motivated contributor made 62,838 edits[2]; and the top 5% of contributors made 44% of the total edits during this time. Such uneven participation has costs: it can lead to a few voices dominating the group and leave the group vulnerable if those few contributors depart the group. Thus, tools that encourage more people to participate may help online communities thrive.

However, motivating contributions to these groups is difficult. As many scholars have observed, online communities and the resources they generate often take the form of a public good, in which all members of the community/public can enjoy the good regardless of their individual levels of contributions (e.g., Ling et al., 2005). Because community members can free ride on others' contributions, people will in general contribute less than would be optimal for the group. Although the critical mass model of collective action (Marwell & Oliver, 1993; Oliver & Marwell, 2001) predicts that a public good can be realized with the contribution of a small number of highly resourceful individuals so long as the provision level of the collective good reaches a level of self-sustainability, involving more contributors can make the group's participation patterns more democratic and robust. To attract contributors the critical mass model (Marwell & Oliver, 1993) maintains that reducing contribution cost is essential.

The cost of contribution to online communities can take multiple forms, including financial cost, emotional cost, and cost in the time and effort in uploading/downloading information. Empirical studies on knowledge management in organizations show that employees were more likely to contribute their expertise to corporate knowledge repositories when contribution did not require too much time or effort (e.g., Yuan et al., 2005). Scholars of online communities have also found that reducing the cost of contribution by improving the design of technologies, e.g. by making it easier to find contributions a person would like to make, could motivate more contributions to a movie website's database(Cosley, Frankowski, Terveen, & Riedl, 2006) or to a discussion group (Ludford, Cosley, Frankowski, & Terveen, 2004).

Following a similar logic, a recommendation tool, WorkFinder was designed and tested in Wikipedia to motivate more contributions to this online information commons. Wikipedia has hundreds of thousands of articles marked as needing improvement (usually lengthening), but no tools to help people find articles they are likely to be able to contribute to—thus, there is a high cost to finding useful contributions to make. Building on the theory of collective action, WorkFinder uses a strategy called "intelligent task routing"[3] to reduce a person's cost of finding articles to work on by recommending articles that are similar to articles that person has edited in the past. Such articles are likely to be close to a person's interests, making it easier for them to contribute.

In this study, we studied the diffusion of WorkFinder in Wikipedia and examined how it influences further contributions. Diffusion of innovation has attracted decades of attention from scholars from diverse disciplines (Burt, 1987; Strang & Soule, 1998; Valente, 1996). However, the difficulties in tracking diffusion processes impose constraints on empirical research. Most studies use retrospective self-report data to examine the diffusion process, and the few studies that collect actual behavior data have sporadic information. For instance, in Coleman et al.'s (1966) study on the diffusion of tetracycline, doctors' prescriptions of the drug were sampled only three consecutive days a month. Errors in recall or gaps in data sampling can add substantial noise to the data, influencing both statistical analysis and conceptual interpretation.

The rise of the Internet has opened up new possibilities for observing diffusion processes. A plethora of digital traces of human online activities can be logged unobtrusively for academic research, giving scholars the opportunity to use objective measures of human behavior over time that are not contaminated with recall biases (Welser, Smith, Gleave, & Fisher., 2008). This may allow researchers to confirm and replicate findings from earlier studies on a much larger scale, as well as to investigate some of the issues that used to be too demanding to study empirically.

When studying the diffusion of WorkFinder in Wikipedia in this study, we were able to obtain a complete record[4] of (a) who has adopted WorkFinder at what time, (b) who has interacted with whom through "user talk pages", Wikipedia's mechanism for users to communicate directly, and (c) who has edited articles suggested by WorkFinder, at what time and with what frequency. We could also find groups of non-adopters, a much overlooked segment in existing diffusion research (Rogers, 2003), as control groups. Finally, because all the data collected has a clear time stamp, the resulting empirical measurements can be arranged along a clear temporal order, which gives us more power to make causal inferences about the relationships among different concepts. Overall, we believe that our research can contribute to diffusion of innovation research from multiple dimensions.

Second, the project improves our understanding of how to motivate contributions to online communities. Motivating contribution to electronic commons is a challenging task because when community members are distributed globally, some conventional incentive strategies, e.g. fostering strong local norms of cooperation (Coleman, 1988), become more difficult to implement. Under this situation, developing new tools to motivate contribution can be a much more effective alternative. Kraut (2003) recommends that information system designers tap into social science theories for design inspirations. WorkFinder, as briefly described above, dovetails with Kraut's call in that its design followed the basic premises of collective action theory (Marwell & Oliver, 1993) and theories of individual motivation to participate in groups (Karau & Williams, 1993), with a goal to involve more people in community development via cost reduction. Through examining WorkFinder's diffusion process, as well as the effect of its adoption, we can better understand how to motivate contributions to online communities.

Using a sample of 6,570 Wikipedia contributors, we explored possible answers to the following questions: (a) what factors influenced adoption of WorkFinder, and (b) whether the adoption of WorkFinder has made a difference in contribution to the community. In the following section of the paper, we will first review related literature about factors that may influence diffusion of innovation and online contributions. We then present an empirical test of the research questions/hypotheses raised. The paper ends with a discussion on substantive implications of our findings, practical implications and directions for future research.

Diffusion of WorkFinder in Wikipedia

Existing studies on diffusion of innovation have identified a long list of factors that can influence the diffusion process. These factors can be classified into roughly two categories: attribute and relational factors (Scott, 1991/2004). Attribute factors focus on individual characteristics, such as innovativeness (i.e., willingness to try out new ideas/products (Rogers, 2003, p. 267-299)), exposure to mass media or metropolitan culture (e.g., Valente, 1996), and so on. Relational factors, focus on structural properties of network relationships, including cohesion, tie homophily, etc. (Burt, 1999; Coleman, Katz, & Menzel, 1957; Rogers, 2003; Strang & Soule, 1998; Valente, 1996). Below, we will review both attribute and relational factors that we think may influence adoption of WorkFinder in Wikipedia.

Individual Attributes

In Wikipedia, we anticipate that highly involved editors are more likely to adopt WorkFinder because these people are more committed to improve the quality of Wikipedia entries. To operationalize level of involvement, we look at two behavioral factors that can be measured from activity logs: *admin status* and *pre-adoption contribution.* Admin status—that is, being listed as an administrator on Wikipedia—is an important indicator of involvement because in Wikipedia only those who have made substantive contributions to Wikipedia over time can earn this status. Obtaining such a status can further motivate contribution for two reasons. First, it is a public acknowledgement of these people's sustained commitment to the community. Second, it gives these contributors additional privileges to help moderate Wikipedia entries and contributors, providing additional incentives to be even more involved with Wikipedia. Given these contributors' intrinsic motivation to improve Wikipedia, we anticipate that they are more likely to adopt tools to cut their cost of contributing. Therefore, it is hypothesized:

Hypothesis 1: Those Wikipedia contributors who have earned admin status before adoption are more likely to adopt WorkFinder than those who did not have such a status.

Admin status is a rather exclusive indicator of involvement in that it is awarded to a relatively small number of contributors. As of Oct. 23, 2008, only 1,612 contributors had admin status out of over 8 million registered users (though most of those are not active contributors). Using admin status alone as an indicator of involvement ignores many contributors who have made substantial contributions to Wikipedia, but who have not earned or sought admin status. Thus, we supplement admin status with another potential behavioral indicator of involvement: amount of pre-adoption contributions. We believe that a positive relationship exists between levels of pre-adoption contribution and the likelihood of adoption because frequent contributors, regardless of their status, share a common motivation with the admin contributors, i.e. to improve the quality of Wikipedia entries. Based on these arguments, it is hypothesized:

Hypothesis 2: High pre-adoption contribution predicts higher likelihood of adoption.

The Influence of Communication Networks

In addition to individual involvement, network connections among contributors can also influence adoption. The turbocharger effect happens when network variables explain additional variance in adoption beyond the direct effects of attribute variables (Rogers, 2003, p. 360). In Wikipedia, contributors communicate with each other through posting on each other's user-talk pages. Social relationships formed over time through such social interactions can significantly influence the likelihood of adoption, causing differences in the time and probability of adoption even though, after an initial test run, WorkFinder was made available to the entire community through a link on a prominent page in Wikipedia.

Social influence through interpersonal exposure. Previous network studies of diffusion found that external influences in the form of subscription to medical journals, cosmopolitan connections, mass media coverage, and so on (DiMaggio & Powell, 1983; Rogers, 2003) can only inform potential adopters of an innovation. It is often “interpersonal influence with friends and neighbors which lead to actual adoption” (Valente, 1996, p. 80). It means that embedded in a network of adopters, a low-involvement contributor may still adopt an innovation following his/her peers after sufficient exposure. Valente maintained that the impact of social influence through direct exposure needs to be accumulated over time. The larger the number of adopters that a focal node has in his/her ego network, the higher the chance of interpersonal exposure. When interpersonal exposures accumulate, potential adopters’ familiarity with the innovation increases (Wejnert, 2002). Over time, the likelihood of adoption grows with reduction in people’s fear and uncertainty about the innovation. When the level of interpersonal exposure exceeds certain threshold, adoption happens. Because WorkFinder writes its suggestions on adopters’ user-talk pages, people who interact with adopters are thus naturally exposed to WorkFinder. Turning these arguments into hypothesis, it is proposed that

Hypothesis 3: High interpersonal exposure predicts higher likelihood of adoption.

Other properties of communication networks can further enhance of the persuasive power of interpersonal exposure. Such network properties include cohesion, tie homophily, and ties to opinion leaders (Rogers, 2003; Strang & Soule, 1998).

Interpersonal cohesion. In the diffusion of innovation research, cohesion refers to the strength of connection between ego and alter (Burt, 1999, p. 39). Cohesion implies a higher level of connectedness with other members in a community, as well as a stronger sense of belonging to the community. Cohesive ties can therefore contribute to higher likelihood of adoption because they enhance the power of social influence during interpersonal exposure (Strang & Soule, 1998). In addition, cohesion implicates a stronger pressure of collective norm (Coleman, 1988). As Wejnert (2002) observed, many cases adoption is a “network-based decision … as pressure toward conformity builds” (p. 306). While people differ in their tendency to confirm to social norms, cohesively tied individuals are likely to mutually influence each other and jointly form a norm that they both are willing to buy into. Based on these arguments, it is hypothesized:

Hypothesis 4: High interpersonal cohesion predicts higher likelihood of adoption.

Tie homophily. Rogers (2003) observed that “interpersonal diffusion networks are most homophilous” (p. 307) because commonalities between adopters and non-adopters increase the power of social influence. Monge and Contractor (2003) summarize two main lines of reasoning that support the theory of homophily, including Byrne’s (1971) similarity-attraction hypothesis and Turner’s (1987) theory of self-categorization. The similarity-attraction hypothesis predicts that people are more likely to interact with those who share similar traits. The theory of self-categorization proposes that people tend to categorize themselves and others in terms of race, gender, age, education, etc. Individuals classified into the same categories perceive themselves more similar to each other. Because interpersonal similarity breeds connections (McPherson, Smith-Lovin, & Cook, 2001, p. 415), social interactions are more likely to happen among similar others. Moreover, because interpersonal similarity and effective communication breed each other (Rogers, 2003, p. 306), homophilous ties become effective means of social influence. The finding of adoption clusters in the diffusion process provides strong evidence for homophilous influence in diffusion. When studying the diffusion of family planning methods in Korean villages, Rogers and Kincard (1981) found “pill villages,” “IUD villages” and “vasectomy villages” where women of the same village tended to adopt the same contraceptive method even though all the different family planning methods were introduced to each village at the same time. In the context of the current research, we anticipate that homophilous ties among contributors who had the same Wikipedian status (i.e. ties among admin contributors and among non-admin contributors) would increase the influence of interpersonal exposure. Over time, peer-to-peer communication via homophilous ties can trigger bandwagon effects among contributors (Abrahamson & Rosenkopf, 1997), and consequently bring about widespread adoption. Summarizing these findings and reasoning, it is hypothesized:

Hypothesis 5: High levels of tie homophily predicts higher likelihood of adoption.

Ties to opinion leaders. While homophilous ties can contribute to adoption, diffusion through homophilous ties tend to be horizontal and confined to people in the same social category (Rogers, 2003). Therefore, ties to people from different social categories is crucial for vertical diffusion throughout the whole system (p. 308). Among different cross-boundary ties, the most crucial ones are those with opinion leaders. Opinion leaders typically occupy higher status in a system, have greater exposure to external information, participate in more social circles, etc. (p.308). Connections with opinion leaders can therefore have a greater influence on adoption because opinions from high-status alters tends to carry more weight. Existing research on health intervention programs has found significant influence of opinion leaders, e.g. breast cancer survivors, on persuading adoption of different prevention practices (Earp et al., 2002).

In Wikipedia, we anticipate that people with admin status can function as opinion leaders. As discussed in Hypothesis 1, because only highly involved contributors can earn admin status, and because the status gives these contributors additional privileges and visibility in Wikipedia, we anticipate that these contributors are more likely than others to adopt WorkFinder. Building on this pre-assumption, we further hypothesize that ties to admin contributors would boost adoption across the whole community. Burt (1999) maintains that opinion leaders are actually information brokers in that they play a key role bringing innovative information from one group to another. Given that opinion leaders are usually better connected (Rogers, 2003), we anticipate that admin contributors can spread WorkFinder beyond the circle of admin contributors and into the community at large. Based on these arguments and reasoning, it is hypothesized that:

Hypothesis 6: More ties to opinion leaders predicts higher likelihood of adoption.

The Impact of Adoption

As discussed earlier, WorkFinder was developed to facilitate contribution to Wikipedia. We were interested in evaluating the effect of adoption, i.e. whether the adoption has resulted in greater contribution to the community. Because the tool was designed to make it easier for contributors to locate which entries need developing or editing, we anticipate that adoption would increase contributions by reducing search cost. It is therefore hypothesized that

Hypothesis 7: Post adoption, adopters of WorkFinder contribute more than non-adopters.

Method

Sample

Our data is based on a set of 6,570 Wikipedia editors. We first selected 2,190 editors who used WorkFinder at least once between March 8, 2006 and March 30, 2007. This is not quite all of the adopters, for technical reasons such as users changing names and occasional errors in the WorkFinder software, but it is the vast majority of adopters. The number of adopters was low compared to the number of contributors to Wikipedia, so we sampled non-adopters following King and Zeng's (2001) recommendation on sampling ratio (between 1:2 and 1:5) for rare event data. We matched two non-adopters with each adopter based on their activity level before adoption and the time they first started editing Wikipedia. These were important factors to control for the following reasons. First edit time was important because even over the course of a year, the behavior and norms of Wikipedia can change significantly. We wanted to ensure that on balance, conditions in Wikipedia were approximately the same for both adopters and non-adopters. We matched users on activity level before adoption to make sure that the adopter and non-adopter groups were comparable. Because the distribution of activity in online communities is highly skewed, a straightforward random sample of non-adopters that did not consider activity would choose low-activity editors almost exclusively.

To do the matching, we looked at every adopter’s first edit time and their adoption time, the date when they first used WorkFinder. For each adopter we chose two non-adopters whose first Wikipedia edit was within a week of the adopter's first edit, and whose number of edits at the corresponding adopter’s adoption time was within 10 percent of the adopter’s number of edits at adoption time. Non-adopters were sampled without replacement (i.e., they would not be selected more than once in the dataset even though one non-adopter may match the profile of two adopters). For a few adopters, we could not find non-adopters who matched the corresponding adopter closely; in these cases, we incrementally widened the first edit date and activity level differences until we could find two matching non-adopters.

Because large samples make it easy to find significant results despite trivial effect sizes, we generated a subsample to evaluate and cross-validate the results. Using the random-generate function in SPSS 16.0, we generated a random subsample of 960 cases for results validation. The subsample was equivalent to 15% of the original sample. All the hypotheses proposed were tested in both the whole sample and the subsample. Despite minor differences in the strengths of relationships, the patterns and directionalities of the tests were consistent across both samples. In the following section, we report results from both samples.

Measurements

Because of the open-source nature of Wikipedia, contributors may be dormant for some time before they become active again. To reduce such noise in data, and exclude activities and social interactions prior to the launch of WorkFinder in Wikipedia, we decided to focus on activity during the 28 days before and after a focal contributor decided to adopt WorkFinder. The 28 day measurement window is intended to capture a snapshot of behavior that is representative of contributors' behavior during a period of recent potential influence.[5] Social network variables were also calculated using the time frame of 28 days before adoption to evaluate the extent of social tie formation through cross-postings on user talk pages. In the following section of the paper, unless specified otherwise, the variables calculated all refer to activities that took place within the 28 days prior to adoption. For non-adopters, the 28 day window was centered on the adoption time of their adopter counterparts to facilitate comparisons.

Individual attribute predictors include several variables that measure an editor's involvement and activity in Wikipedia. *Admin status* was a dummy variable with 1 representing "have obtained administrator status in Wikipedia by the time of adoption," and 0 representing "have not." *Pre-adoption contribution* measures the total number of edits that a contributor has made during 28 days before adoption. *Total pre-adoption contribution* measures the total number of edits that a contributor has made since joining Wikipedia. *Total post-adoption contribution* measures the total number of edits that a contributor has made since adoption.

Network variables were generated by tracing Wikipedia contributors' posts on each other's user-talk pages. *Interpersonal exposure to innovation* was calculated by counting the number of adopters of WorkFinder that a contributor had direct contact with during 28 days prior to adoption. *Interpersonal cohesion* was measured by counting the number of reciprocal ties in a contributor's ego network. *Tie homophily* was calculated using the percentage of a contributor's same-status ties (admin vs. non-admin). *Finally, ties to opinion leaders* were measured by counting the number of ties that an ego had with contributors who had admin status.

Control variable. In addition to the research variables, *number of active months* was calculated by counting the number of calendar months in which a contributor has made at least one edit before April, 2007. We believe that this measurement of tenure with Wikipedia is a stronger control variable than a mere count of the number of months since a person has joined because given the voluntary nature of this open content community, it is common that contributors remain dormant for some time before turning into active contributors again.

Analysis Method

To address the hypotheses around diffusion, we used logistic regression models because the dependent variable had only two response categories, with 1 = Adopted and 0 = Didn't Adopt. In reporting the results, both the partial regression coefficients (B) and their corresponding odds ratios are included because the odds ratios reflect more directly the likelihood of adoption vs. non-adoption as predicted by a particular independent variable. When the odds ratio is greater than 1, the independent variable can predict the likelihood of adoption better than chance alone (50%). The corresponding *p* value of the Wald test reveals whether the improvement in prediction is statistically significant. The multiple R^2, which is calculated on the basis of deviance scores, measures variance explained in the dependent variable. Differences in deviance scores between nested models follow a χ^2 distribution. A significant change in χ^2 implies significant improvement in fit of the regression equation with the data.

To address the question about the impact of adoption on contribution to Wikipedia, we conducted independent sample t-tests to compare the adopter and the non-adopter groups in their level of contribution to Wikipedia community both *before* and *after* adoption. Significant changes in contribution levels before and after adoption between adopters and non-adopters reveal the impact of adoption.

Data Preparation

Prior to conducting a series of tests of the hypotheses, a number of variables were recoded to reduce the skewness in distribution. As described earlier, in most online communities, a small group of people tend to contribute a disproportionately high share of content, while most members are much less active, causing extremely high skewness in data distribution. Although logistic regression does not require normal distribution of variables (Tabachnick & Fidell, 1996, p.575), departure from normality does influence the robustness of estimates. Among the three commonly-used data transformation methods to improve normality of distribution, i.e. square root, log, and inverse, Osborne (2002) maintained that inverse was most effective in transforming extremely skewed distributions. Following his suggestion, we first took the inverse of a variable. We then multiplied the inversed value by -1 because the first step of transformation reversed the rank order of variables. Finally, a constant was added to the distribution to bring the minimum value of these variables above 1.0. After these linear transformations, the skewness levels of these recoded variables were reduced significantly[6]. The descriptive statistics and zero order correlations among the transformed variables are reported in Table 1. The relationship between admin status and other variables were bi-serial correlation coefficients because admin status was a dummy variable; the rest were Pearson correlation coefficients.

>>>>>Insert Table 1 about here <<<<<

Results

Hypotheses 1 to 6 examined what factors influenced adoption. Since they shared the same dependent variable, the research variables were entered into the logistic regression model in steps; the models are shown in Table 2. Model 1 contained only the control variable, number of active months. When number of active months was the only predictor variable in the analysis, it had significant influence on likelihood of adoption in both the whole sample (B=.02, odds ratio of =1.02, p<.05), and the subsample (B=.02, odds ratio of =1.02, p<.05). The changes in the deviance scores from the null, baseline model showed that the improvement in model fit was significant in both the whole sample (χ^2=51.70, df=1, p <.05), and the subsample (χ^2=5.48, df=1, p <.05). However, this model explained only 1% of variance in the likelihood of adoption.

>>>>>Insert Table 2 about here <<<<<

Model 2 added the attribute variables of admin status and total activity. The control variable, number of active months, became non-significant when the attribute variables were added. Counter to Hypothesis 1, admin status was not a significant predictor of likelihood of adoption in either the whole sample (B=.05, odds ratio=1.05, p >.05) or the subsample (B=.37, odds ratio=1.45, p >.05). That is, although the odds ratios showed that those with admin status had higher likelihood of adopting WorkFinder, the increased likelihood was not statistically significant. Consistent with Hypothesis 2, pre-adoption contribution was a significant predictor of adoption in both the whole sample (B=2.42, odds ratio=11.23, p <.05) and the subsample (B=2.48, odds ratio=11.89, p <.05). That is, heavy contributors had much higher likelihood of adoption. The two attribute variables explained 14% of additional variance in the dependent variable compared to Model 1. Changes in deviance scores proved that the improvement in model fit between the control-variable-only model and the current model was significant in both the whole sample (χ^2=709.64, df=2, *p* <.05), and the subsample (χ^2=105.56, df=2, *p* <.05).

Model 3 added four network variables. Consistent with Hypothesis 3, interpersonal exposure was a significant predictor of likelihood of adoption in both the whole sample (B=1.12, odds ratio=3.08, *p* <.05) and the subsample (B=1.61, odds ratio=5.01, *p* <.05). Also consistent with Hypothesis 4, cohesion through reciprocal ties was a significant predictor of likelihood of adoption in both the whole sample (B=.66, odds ratio=1.93, *p* <.05) and the subsample (B=.90, odds ratio=2.45, *p* <.05). Supporting Hypothesis 5, tie homophily was a significant predictor of likelihood of adoption in both the whole sample (B=.68, odds ratio=1.97 *p* <.05) and the subsample (B=.85, odds ratio=2.34, *p* <.05). However, counter to Hypothesis 6, ties to opinion leaders did not increase likelihood of adoption in either the whole sample (B=-.13, odds ratio=.88, *p* >.05) or the subsample (B=-.79, odds ratio=.46, *p* >.05). Model 3 explains 22% variance in the likelihood of adoption in both samples. The amount of change in deviance score showed that the observed improvement in model fit from Model 2 was statistically significant in both the whole sample (χ^2=359.46, df=4, *p* <.05), and the subsample (χ^2=71.32, df=4, *p* <.05).

The second research question explores how adoption affects contribution. Hypothesis 7 predicts that adopters would contribute more in the future than nonadopters. To test this, we conducted independent sample *t*-tests to compare the mean level of contributions between adopter and non-adopter groups both *before* and *after* adoption. Levene's test showed that equal variance between the two groups in the whole sample could not be assumed (F(1, 6,568)=6.37, p<.05). The corresponding *t* value was statistically significant, *t* (4,776) =-2.34, p < .05, indicating that before adoption, adopters (M=926.52, SD=1875.71) on average contributed less to the community than the non-adopters (M=1,049.55, SD=2,067.72). However, the contribution pattern was reversed after adopters adopted WorkFinder. An independent sample *t*-test on post-adoption contribution found significant difference in means, *t* (4,317) = 4.99, p<.01 between adopters (M=788.86, SD=1,741.31) and non-adopters (M=562.82, SD=1,713.83), indicating that adopters contributed significantly more than the non-adopters after they adopted WorkFinder.

The same independent sample *t*-tests were also conducted with the randomly generated subsample. The tests produced similar results. Before adoption, adopters (M=769.65, SD= 1,613.94) and non-adopters (M=910.51, SD= 1,677.96) did not differ in their levels of contribution to the community (equal variance of two groups can be assumed (F(1, 958)=2.41, p>.05), *t* (958) = -1.24, p>.05). However, after adoption, the adopters (M=787.36 SD=1805.67) contributed significantly more to the community than non-adopters (M=479.75, SD=1,806.83), *t* (638) = 2.49, p<.01 (equal variance between the two groups could not be assumed (F(1, 958)=4.30, p<.05)). Taken together, the results based on both the whole sample and the subsample supported Hypothesis 7, suggesting that adoption of WorkFinder enabled adopters to make significantly more contributions to the community even when non-adopters contributed equally or more than adopters prior to adoption.

Discussion

Just as cities need a visitor's bureau and a sanitation department, online communities need to welcome newcomers, create valuable content, and police problems. A key problem for these communities is motivating their members to perform this kind of maintenance activity. The critical mass model of public good suggests that reducing the cost of contribution can motivate contributions (Marwell & Oliver, 1993). Building on this proposition, WorkFinder was developed to make it easier for Wikipedia editors to locate entries that match their abilities to contribute. In this study, we examined how WorkFinder diffused in Wikipedia and whether the adoption of WorkFinder has made any impact on people's contributions. The research aims at contributing to our understanding of both diffusion of innovation processes and the development of online communities. The results confirmed some of the insights from earlier diffusion studies. At the same time, they also point out some areas of research worthy of further investigation.

Contributions

From hybrid corn to poison pills (Strang & Soule, 1998), diffusion of innovation as a research paradigm has shown wide applicability across multiple disciplines. The appeals and practical implications of the paradigm are timeless because the success of an innovation would be incomplete without successful diffusion. However, some research issues remain inadequately explored even after decades of research on the topic. Rogers observed (Rogers, 2003; 1971) a pro-innovation bias in diffusion of innovation research in that researchers tend to focus more on adopters of successfully diffused innovations. As a result, we have very limited knowledge about non-adopters and unsuccessfully diffused innovations (Rogers, 2003). While adopters and non-adopters may share many commonalities, it is premature to accept this assumption without empirical evidence—but such evidence is hard to collect. Using retrospective self-reported data to study only adopters of successfully diffused innovation is unavoidable in many situations because diffusion is a process that has an unknown prospective duration of diffusion time. It is therefore hard for researchers to predict in advance the right time to track the actual behavior of every person in the population when working with limited resources. Thus, empirical difficulties in tracking diffusion processes can potentially impose major conceptual constraints on what questions can be asked and/or addressed.

Reaping the benefit of a growing body of digital traces of people's natural behavior on the Internet, we studied how individual and network factors influenced diffusion of an intelligent task recommendation tool, WorkFinder, in Wikipedia. Instead of focusing on adopters only, we first randomly sampled 2,190 adopters, obtained the public records of their contribution to Wikipedia, and then matched each adopter in our sample with two non-adopters by their starting time and their overall activity level prior to adoption. The inclusion of these non-adopters helped us gain a better understanding about the key factors that influence adoption in the whole population. Moreover, matching adopters and non-adopters on their pre-adoption activity levels and time of joining the community provides an even more conservative test of the research hypotheses because the sampling procedure made non-adopters in our comparison group resemble adopters more closely than a randomly selected non-adopter in multiple dimensions, including network connections, likelihood of adoption, and involvement with Wikipedia.

About individual attribute variables, we anticipated that those contributors with admin status would show higher likelihood of adoption, given they are more likely to have sustained interest in improving Wikipedia. This was not the case; though admins had a higher likelihood of adoption, the difference was not statistically significant. On the other hand, the hypothesis that high levels of contribution, as a second indicator of level of involvement with Wikipedia, would predict adoption was strongly supported in both the whole sample and the subsample. Because adopters and non-adopters were matched on overall pre-adoption activity levels[7], and the variable measuring contribution in the analysis mainly focused on total contributions within 28 days prior to adoption, the result suggests that the intensity of involvement right before adoption is a more important predictor of adoption than overall activity.

Consistent with our hypotheses, the network factors of interpersonal exposure, cohesion through interpersonal ties and tie homophily were all found to be significant predicators of higher likelihood of adoption in both the whole sample and the subsample. In combination, the results suggest that an important tool for increasing the likelihood of innovations spreading through a community is building stronger ties between members, giving members more opportunities to receive direct exposures to an innovation (Hypothesis 3) via cohesive ties (Hypothesis 4) among contributors of similar characteristics (Hypothesis 5).

Overall, our hypotheses focused on the importance of cohesion for diffusion. Burt (1987, 1999) on the other hand, provided some strong arguments about the importance of structural equivalence for diffusion in a variety of off-line contexts. Because both online ego networks and complete social networks are likely to have fuzzy, ephemeral boundaries, evaluation of structural equivalence among a random sample of 6,570 contributors among 8 million may be difficult. As Scott pointed out (1991/2004), structural equivalence mainly focuses on comparisons among different nodes within the same finite network. When the network is huge and its boundary is unclear, the influence of structural equivalence on diffusion is murkier. We do not want to rule out completely the importance of structural equivalence for diffusion in an online environment. On the other hand, it may be a reality that people find it harder to identify their structurally equivalent counterparts online unless they are connected to their counterparts via cohesive ties at the same time. As a result, it may be very challenging to prove that the drive for competition would motivate adoption among structurally equivalent actors online, as Burt has described.

Finally, we hypothesized that the admin contributors, given their status, as well as the visibility of their status, would function as both opinion leaders and information brokers (Burt, 1999) to help spread the innovation across the whole community. Counter to our prediction, ties to admin contributors did not increase likelihood of adoption. Since admin status did not predict adoption (Hypothesis 1), this finding became less surprising. The lack of wide support from across the high-status members of the community may explain the relatively low level of adoption of WorkFinder, despite the observed significant increase in contribution by adopters (Hypothesis 7) post adoption. That is, without opinion leaders leading the way, WorkFinder did not diffuse widely through the community, which adds support to earlier studies' findings about the important role of opinion leaders in diffusion processes (Earp et al., 2002; Rogers, 2003).

Empirically, our study contributes to a growing trend in social science research to use large records of people's activities online to study human behavior (Crandall, Cosley, Huttenlocher, Kleinberg, & Suri, 2008; T. Turner, Smith, Fisher, & Welser, 2005). With the use of these large behavioral datasets, however, come a different set of methodological challenges. Behavioral data is often non-normally distributed, affecting both activity measurement and sampling strategies (Welser et al., 2008). Naïve sampling methods may over- or under-represent the null case in logistic regression models, as King & Zeng (2001) point out. We point out here that naïve random sampling may also lead to inappropriate comparison groups without careful consideration of the variables on which the comparison should be matched, in terms of the dynamics of the community being studied. Wikipedia has grown exponentially for years, and activity is roughly exponentially distributed; our sampling strategies needed to take this into account. Our choice of 28-day windows was driven by the need to walk a fine line between failing to detect individual-level change (longer periods) and being driven by the natural, intermittent nature of activity in online contexts, where a sudden emergency such as a paper deadline might curtail a person's activity in the online community. Our windows are also participant-specific, centered around events of particular interest, which is potentially more informative than standard strategies that use calendar time to divide time series data into periods. Working with computer programmers and large datasets allowed us to do complex, nuanced sampling, and this is likely to become more common and more important in social science research going forward.

Limitations

One major limitation of the current research is that we used only archival data to study Wikipedia contributors' online activities. While the logged data provided very precise, objective measures of a number of key variables related to diffusion, network relationships, social interactions, etc., using archival data alone has some limitations. For instance, archival data provides behavioral indicators around contribution, but no psychological measures of what motivates contribution to the Wikipedia community. In addition, because we relied on Wikipedia archive data exclusively, we did not have information about the frequency or reciprocity of communication outside of Wikipedia, either in other online channels or in offline contexts. Wikipedians, especially admins and other committed members, do have offline meetings and dedicated communication channels outside of Wikipedia; how much these affect their behavior is impossible to tell from the available archives. Richer data, including survey and interview responses along with information from other communication media, can reveal much more information about the differences between active and not-so-active contributors in their interests, motivations, experiences with Wikipedia platforms, emotional involvement with the community, and so on. Findings from mixed methods can both give a richer picture of what is happening and perhaps, as in Crandall et al. (2008), help with the quantitative modeling work directly.

Direction for Future Research

This paper has only begun to reveal the potential of studying diffusion processes using log files of people's behaviors online. We see several directions for further research that advances our understanding of diffusion. First, as Rogers (2003) pointed out, while it has been widely studied what factors influence one-time adoption, little work has examined what factors will influence continued usage—or discontinuance—of an innovation (p. 110). Rogers and Shoemaker (1971) first observed this pro-innovation bias in diffusion of innovation research around four decades ago. Still not much has been done to tackle the issue (Rogers, 2003). As a result, we know much more about adoption and use than about continued use and discontinuance (p. 111). Yet continued usage of innovation is important because it is a more powerful indicator of success. We believe that the difficulties in collecting empirical data contribute to the focus on one-time adoption. In the current research, we were able to collect data on whether adopters of WorkFinder used the tool repeatedly to locate entries that need their work. While adopters tended to contribute more to Wikipedia post-adoption (as shown in the independent sample t-test reported earlier), the diffusion of WorkFinder was not as successful as we have anticipated in either the scope or the depth of adoption. WorkFinder adopters account for well under 1% of the whole population of registered Wikipedia users. Further, relatively few adopters, about 15%, use WorkFinder repeatedly. This number has grown since WorkFinder offered a subscription option that allows adopters to receive suggestions on a repeated basis, but it is still small. It would be interesting to conduct a follow up study to find out why non-adopters fail to use WorkFinder. It would also be interesting to survey or interview those adopters who did not use the tool repeatedly to find out why they have stopped using the tool. For our specific context, insights from these studies might lead to a better design for WorkFinder; more generally, they may point to factors that could inform future diffusion research. For example, our prior experience developing innovations leads us to believe continued use heavily depends on initial experiences with an innovation, but empirically showing how important this is compared to attribute and relational factors could inform both design and research around innovation.

Second, in post-hoc exploratory analysis, we found that structural properties of social network, including interpersonal cohesion, ties to opinion leaders, etc. did not make a huge difference in the level of post-adoption contributions to Wikipedia. One primary reason is that the primary goal of WorkFinder is to motivate contributions from each individual editor. While new network ties may develop among editors who receive recommendations to edit similar articles, the tool did not have a component designated for fostering the development of social ties among contributors. In future research, it would be interesting to explore whether the addition of such a component would help foster a greater sense of belonging to a community, and thereafter motivate more contributions to Wikipedia. It will also be interesting to compare whether a tool with a networking component diffuses differently from a tool that focuses exclusively on facilitating individual contributions. Other datasets might be especially appropriate for this kind of analysis, such as data drawn from explicit social networking sites like Facebook.

Practical Implications

While the adoption of WorkFinder in Wikipedia is not as widespread as we would have liked, the results showed that the design and implementation of intelligent-routing task recommendation systems in online communities can significantly increase contribution. The research shows the exciting promise of using technology as interventions to boost online community involvement. Articles WorkFinder suggests are edited about four times as often as randomly chosen articles; it has received dozens of positive comments and several awards; and other wikis are interested in WorkFinder. Intelligent task routing also works in other communities such as MovieLens (http://www.movielens.org/) for similar public goods such as databases of movie information. The combination of its success in multiple communities, the theories of motivation and collective goods that underlie the idea, and the fact that simple recommendation strategies are effective all suggest that intelligent task routing is a valuable, general idea that could help make many online communities better.

This work also points to important design considerations for increasing the likelihood that their idea or innovation will propagate through an online community. The fact that people were exposed to WorkFinder by seeing it on other users' talk pages was a fortunate side effect of the way it presented its suggestions, rather than a planned strategy for increasing its diffusion. Designing the interface and appearance of an innovation so that it naturally taps into the power of network variables for supporting diffusion is an important strategy, especially in the rapidly growing world of online social systems, from Wikipedia to blogs to Facebook. Badly done, this can hurt innovation—for instance, Facebook applications sometimes encourage people to send unsolicited invitations to everyone in their social network. Most of these ties are low cohesion and little homophily of interests, so in light of our findings it is not surprising that this practice often fails to spread the application and in fact often creates a backlash against it. Our work suggests that using only strong links and links between people who are fairly similar may be a more effective strategy for effectively diffusing an innovation

Conclusion

Overall, this work confirms, extends, and informs both intuitions about influence in social networks (e.g. Gladwell's *The Tipping Point*) and studies of actual diffusion in networks. It shows the promise of carrying out diffusion studies in large-scale online interaction log data, demonstrates important factors in the diffusion of WorkFinder through the Wikipedia community, suggests that such tools will work to increase contributions to these communities, points out important methodological issues in doing this kind of work on large datasets, and shows potential design opportunities for future innovators based on its results.

Endnotes:

[1] Scholars disagree in their definitions about what constitute online communities. Following Preece and Maloney-Krichmar (2005), we think it is more important to “concentrate on more substantive issues such as how communities are created, evolve or cease to exist online” (p. 2) than to clean the “fuzzy boundaries” of a concept “that is more appropriately defined by membership” (p. 2).

[2] It is likely this editor used a tool, in Wikipedia parlance called a “bot”, that helps editors make a large batch of related edits (such as correcting a particular misspelling in a number of articles).

[3] WorkFinder uses several technologies for recommending articles, similar to technologies that help companies like Amazon recommend books and music. The related algorithms and other technical details of the tool are described in detail in _____(the specific reference is removed here for blind review.

[4] A data file was generated partially from the public record of who has talked to whom, who has edited which articles, etc. that all Wikipedia editors can see, and partially from the log file of WorkFinder usage activities.

[5] All such windows entail trade-offs between windows that are too long (thus masking long term individual level change) and windows that are too short (and thus are more strongly affected by the intermittent nature of participation).

[6] We also tried the log 10 based transformation. Although the results were consistent with the inverse transformation, the skewness level of the variables after log transformation was worse than those after inverse transformation.

[7] The independent sample *t*-test conducted to test Hypothesis 7 confirmed the success of matching because adopters and non-adopters were not statistically different from each other in their overall levels of contribution prior to adoption.

References:

Abrahamson, E., & Rosenkopf, L. (1997). Soical network effects on the extent of innovation diffusion: A computer simulation. *Organization Science, 8*(3), 289-309.

Burt, R. S. (1987). Social contagion and innovation: Cohension versus structural equivalence. *Amercian Journal of Sociology, 92*(May), 1281-1335.

Burt, R. S. (1999). The social capital of opinion leaders. *Annals of the American Academy of Political and Social Science, 566*, 37-54.

Byrne, D. E. (1971). *The attraction paradigm*. New York: Academic Press.

Coleman, J. (1988). Social capital in the creation of human-capital. *American Journal of Sociology, 94*, S95-S120.

Coleman, J., Katz, E., & Menzel, H. (1957). The diffusion of an innovation among physicians. *Sociometry, 20*(4), 253-270.

Coleman, J., Katz, E., & Menzel, H. (1966). *Medical innovation: A diffusion study*. New Yoirk: Bobbs Merrill.

Cosley, D., Frankowski, D., Terveen, L., & Riedl, J. (2006). Using intelligent task routing and contribution review to help communities build artifacts of lasting value. In R. Grinter, T. Rodden, P. Aoki, E. Cutrell, R. Jeffries & G. Olson (Eds.), *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1037-1046). Montréal, Québec, Canada: ACM New York, NY, USA.

Crandall, D., Cosley, D., Huttenlocher, D., Kleinberg, J., & Suri, S. (2008). Feedback effects between similarity and social influence in online communities. In

Proceeding of the 14th ACM SIGKDD international Conference on Knowledge Discovery and Data Mining (pp. 160-168). New York, NY: ACM.

Cummings, J. N., Butler, B., & Kraut, R. (2002). The quality of online social relationships. *Communications of the ACM, 45*(7), 103-108.

DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review, 48*(2), 147-160.

Earp, J. A., Eng, E., O'Malley, M. S., Altpeter, M., Rauscher, G., Mayne, L., et al. (2002). Increasing use of mammography among older, rural African American women: results from a community trial. *American Journal of Public Health, 92*(4), 646-654.

Karau, S. J., & Williams, K. D. (1993). Social loafing: A meta-analytic review and theoretical integration. *Journal of Personality and Social Psychology, 65*(4), 681-706.

King, G., & Zeng, L. (2001). Logistic Regression in rare event data. *Political Analysis, 9*(2), 137-163.

Kraut, R. E. (2003). Applying social psychological theory to the problems of group work. In J. Carroll (Ed.), *HCI models, theories and frameworks: Toward a multidisciplinary science* (pp. 325-356). New York: Morgan Kaufman.

Lakhani, K. R., & Hippel, E. V. (2003). How open source software works: "Free" user to user assistance. *Research Policy (Special Issue), 32*(6), 923-943.

Ling, K., Beenen, G., Ludford, P., Wang, X. Q., Chang, K., Li, X., et al. (2005). Using social psychology to motivate contributions to online communities. *Journal, 10*(4). Retrieved from http://jcmc.indiana.edu/vol10/issue4/ling.html

Ludford, P. J., Cosley, D., Frankowski, D., & Terveen, L. (2004). Think different: increasing online community participation using uniqueness and group dissimilarity. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 631-638). New York, NY: ACM.

Marwell, G., & Oliver, P. (1993). *The critical mass in collective action: A micro-social theory*. New York City, NY: Cambridge University Press.

McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology, 27*, 415-444.

Mockus, A., Fielding, R. T., & Andersen, H. (2002). Two case studies of open source software development: Apache and Mozilla. *ACM Transactions on Software Engineering and Methodology, 11*(3), 309-346.

Monge, P. R., & Contractor, N. (2003). *Theories of communication networks*: Oxford University Press.

Oliver, P., & Marwell, G. (2001). Whatever happened to critical mass theory? A retrospective and assessment. *Sociological Theory, 19*(3), 292-311.

Osborne, J. (2002). Notes on the use of data transformations. *Journal, 8*(6). Retrieved from http://PAREonline.net/getvn.asp?v=8&n=6

Preece, J., & Maloney-Krichmar. (2005). Online communities: Design, Theory and Practice. *Journal of Computer-Mediated Communication, 10*(4).

Rogers, E. M. (2003). *Diffusion of innovation* (5th ed.). New York: Free Press.

Rogers, E. M., & Kincaid, D. L. (1981). *Communication Networks: Toward a New Paradigm for Research*: Free Press.

Rogers, E. M., & Shoemaker, F. E. (1971). *Communication of innovations: A cross-cultural approach*. New York: Free Press.

Scott, J. (1991/2004). *Socia.l network analysis*. Thousand Oaks, CA: Sage.

Strang, D., & Soule, S. A. (1998). Diffusion in organizations and social movements: From hybrid corn to poison pills. *Annual Review of Sociology, 24*, 265-290.

Tabachnick, B. G., & Fidell, L. S. (1996). *Using multivariate statistics* (3rd ed.). New York: HarperCollins Publisher

Turner, J. C. (1987). *Rediscovering the social group: A self-categorization theory*. Oxford, UK: Basil Blackwell.

Turner, T., Smith, M. A., Fisher, D., & Welser, H. T. (2005). Picturing usenet: Mapping computer-mediated collective ation. *Journal, 10*(4). Retrieved from http://jcmc.indiana.edu/vol10/issue4/turner.html

Valente, T. W. (1996). Social network thresholds in the diffusion of innovations. *Social Networks*, *18*(1), 69-89.

Wejnert, B. (2002). Integrating models of diffusion of innovations: A conceptual framework. *Annual review of sociology, 28*, 297-326.

Welser, H. T., Smith, M. A., Gleave, E., & Fisher., D. (2008). Distilling digital traces: Computational social science approaches to studying the Internet. In N. Fielding, R. L. Lee & G. Grant (Eds.), *Handbook of online research methods* (pp. 116-140). London: Sage Publications.

Yuan, Y. C., Fulk, J., Shumate, M., Monge, P., Bryant, J. A., & Matsaganis, M. (2005). Individual participation in organizational information commons: The impact of team level social influence and technology-specific competence. *Human Communication Research, 31*(2), 212-240.

Table 1

Descriptive Statistics and Zero-Order Correlations

	1	2	3	4	5	6	7	8	9	10	Mean	Standard Deviation
1. Active months		.37*	.06	.50*	.16*	.05	.10*	.07*	.14*	.08*	13.21	9.35
2. Admin status	.35*		.05	.42*	.09*	.25*	.21*	-.07*	.29*	.06	.03	.18
3. Pre-adoption contribution (28 days before adoption)	.14*	.10*		.33*	.46*	.28*	.35*	.28*	.40*	.10*	107.03	200.83
4. Total contribution before adoption	.47*	.43*	.46*		.26*	.10*	.15*	.07*	.20*	-.04	863.56	1657.38
5. Total contribution after adoption	.23*	.13*	.56*	.42*		.15*	.20*	.15*	.22*	.08*	582.29	1811.32
6 .Interpersonal exposure	.06*	.12*	.19*	.11*	.10*		.79*	.16*	.78*	.24*	.66	2.10
7. Cohesion	.15*	.24*	.40*	25*	.26*	.52*		.19*	.82*	.23*	1.74	5.32
8. Tie homophily	.13*	-.10*	.27*	.09*	.18*	.14*	.21*		.14*	.28*	.45	.44
9. Ties to opinion leaders	.16*	.30*	.33*	.24*	.23*	.53*	.80*	.12*		.21*	1.40	3.73
10. Adoption	.09*	.02*	.10*	-.03*	.06*	.13*	.22*	.27*	.18*		.33	.47
Mean	13.65	.04	112.75	1008.54	638.17	.69	1.82	.45	1.42	.33		
Standard Deviation	9.72	.19	239.22	2006.45	1726.20	2.87	4.91	.44	4.08	.47		

*p<.05

Note: Lower triangular of the matrix reports the descriptive statistics and zero-order correlation using the whole sample (N=6570); upper triangular of the matrix reports the descriptive statistics and zero-order correlation using the random sample (N=960).

Table 2

Results of Logistic Regression Analysis

Variable	Model 1		Model 2		Model 3	
	Whole Sample	Subsample	Whole Sample	Subsample	Whole Sample	Subsample
Number of active months	1.02**	1.02*	1.00	1.01	1.00	1.00
Admin status			1.05	1.45	.77	1.05
Pre-adoption contribution			11.23**	11.89**	5.14**	5.42**
Interpersonal exposure					3.08**	5.01*
Cohesion					1.93**	2.45*
Tie homophily					1.97**	2.34**
Ties to opinion leaders					.88	.46
R^2	0.01	0.01	0.15	0.15	0.22	0.24

** p<.01 *p<.05
Whole sample (n=6,570)
Subsample (n=960)