

Membership Turnover and Collaboration Success in Online Communities:

Explaining Rises and Falls from Grace in Wikipedia

Sam Ransbotham
Carroll School of Management
Boston College
140 Commonwealth Ave
Chestnut Hill, MA 02467
sam.ransbotham@bc.edu

Gerald C. Kane
Carroll School of Management
Boston College
140 Commonwealth Ave
Chestnut Hill, MA 02467
gerald.kane@bc.edu

forthcoming in MIS Quarterly

Abstract

Firms increasingly turn to online communities to create valuable information. These communities are empowered by new information technology-enabled collaborative tools, tools such as blogs, wikis, and social networks. Collaboration on these platforms is characterized by considerable membership turnover, which could have significant effects on collaborative outcomes. We hypothesize that membership retention relates in a curvilinear fashion to effective collaboration: positively up to a threshold and negatively thereafter. The longitudinal history of 2,065 featured articles on Wikipedia offers support for this hypotheses: Contributions from a mixture of new and experienced participants both increases the likelihood that an article will be promoted to featured article status and decreases the risk it will be demoted after having been promoted. These findings imply that, contrary to many of the assumptions in previous research, participant retention does not have a strictly positive effect on emerging collaborative environments. Further analysis of our data provides empirical evidence that knowledge creation and knowledge retention are actually distinct phases of community-based peer production, and that communities may on average experience more turnover than ideal during the knowledge retention phase.

Keywords: *online communities, collaboration, longitudinal study, membership turnover, information generation, information retention.*

INTRODUCTION

Due in no small part to the emergence of a new class of Internet-based collaborative tools, commonly known as Web 2.0 or social media, companies are increasingly turning to online communities as sources of valuable information. New companies such as Threadless (Lakhani and Kanji 2008) and Communispace (Li and Bernoff 2008) rely on input from communities to offer entirely new business models. Traditional companies such as Dell (DiGangi et al. 2010) and Starbucks (Gallaughier and Ransbotham 2010) also use social media to cultivate online communities and thereby solicit and evaluate product development ideas from customers. Despite anecdotal success, many social media communities fail to generate any worthwhile information. One consulting firm thus estimates that the majority of *Fortune* 1000 firms will experiment with social media communities, but more than half of their efforts will fail to generate desired outcomes (Sarner 2008). Some failures may be due to technical reasons, but the majority will fail because they cannot generate effective collaborative processes among participants. For example, when the *Los Angeles Times* attempted to use a social media platform to capture opinions about the involvement of the U.S. military in Iraq, the collaboration devolved quickly, as participants on one side of the debate simply deleted and replaced contributions from the other side (Wagner and Majchrzak 2006).

Part of the reason for such failures may reflect a key aspect of collaboration in online communities, namely, the high levels of membership turnover (Faraj et al. forthcoming; Kane et al. 2009; Oh and Jeon 2007). Without committing to any tasks, projects, or conversations, each participant is free to come and go. This turnover creates a continuously changing environment in which active participants rarely remain the same over time. Some participate for mere minutes; others remain for longer. Some participants make a single contribution, whereas others offer

substantial contributions that require their considerable effort and energy. The online community literature generally contends that turnover is detrimental to effective collaboration, and the ability to attract and retain members represents a key metric of success (e.g., Arguello et al. 2006; Butler 2001; Lazar and Preece 2002; Ma and Agarwal 2007). Every time a participant leaves the community, he or she takes not only unique knowledge and insight but also the experience that person has gained through participation. It has been posited, therefore, that these departures diminish the resources available to the community and may threaten its very sustainability. Only if participants remain in the community will they gain experience and insight that can be applied to improve individual and collective collaboration.

Yet some capabilities of social media platforms could help mitigate the negative effects of turnover. For example, social media platforms typically preserve all previous contributions by past members in an organized, searchable format, such that future community members can modify and adapt their contributions as needed (Kane and Fichman 2009; Wagner and Majchrzak 2006). The platforms often also provide separate forums for discussing collaboration issues that encourage consensus, and these collaborative decisions can be preserved for future members. Some research suggests that these information technology (IT)-enabled features help mitigate the negative effects of turnover (Kane and Alavi 2007). If the platform can effectively retain the contributions by its members, some amount of turnover might benefit collaboration by allowing new participants to offer insights and knowledge that the community previously did not possess. It may be that, by retaining the entire collaborative history, the platform creates past experience for the community that its individual members lack.

This paper investigates how membership turnover affects collaborative outcomes in social media communities by examining the entire collaborative history of 2,065 “featured”

articles on Wikipedia. *Featured articles* are those that Wikipedia recognizes as the best exemplars of the type of information Wikipedia seeks to generate. We hypothesize that membership retention relates in a curvilinear fashion to effective collaboration (positively up to a threshold and negatively thereafter). Kane et al. (2009) suggests two stages of collaboration are critical for examining collaboration in social media communities, namely, the *creation stage* when information is developed and shaped, and the *retention stage* when the created information gets preserved and refined through ongoing collaboration. Collaboration and turnover continue after the community has collaborated successfully, and social media communities must both create and retain knowledge. We find support for the curvilinear relationship between membership turnover and performance in both stages, and our control variables support the distinction of promotion and demotion as different collaborative stages. We also find that communities on average experience more turnover during the knowledge retention phase than would be optimal for effective collaboration.

Our analysis also finds empirical evidence that knowledge creation and knowledge retention are distinct phases of community-based peer production. The collaboration that occurs during the knowledge creation phase has little effect on the effectiveness of later knowledge retention. Furthermore, the factors associated with effective knowledge retention are also different from those associated with effective knowledge creation leading to important implications for practice and theory.

Three Views of Turnover in Organizations

Considerable literature investigates the effect of turnover on performance in traditional organizations, but perspectives on this relationship vary widely. The most common view holds that turnover relates negatively to performance (Huselid 1995; Ton and Huckman 2008). Why

would this be so? When people leave, the organization must expend resources to recruit and train new employees to replace them (Darmon 1990; Hom and Griffeth 1995; Staw 1980). Departing employees can take unique experience and knowledge with them (Argote and Epple 1990; Becker 1962; Carley 1992; Nelson and Winter 1982) or their leaving may disrupt the social networks or work environment of those who remain (Dess and Shaw 2001; Leana and Van Buren 1999). Whether turnover incurs replacement costs, disrupts the work environment, or weakens the knowledge resources of the organization, the conventional view is that turnover harms organizational performance (Glebbeck and Bax 2004; Huselid 1995).

Despite the dominance of this view, it is by no means the only perspective. Another argument suggests that turnover in certain situations may benefit organizations because those who leave often are those most dissatisfied with the current organization, such that those who remain behind enjoy better working conditions and performance (Krackhardt and Porter 1985). Furthermore, IT-based platforms now allow organizations to collect and store employee knowledge, so an employee's worth to the organization actually declines once their knowledge has been stored in a knowledge repository (e.g., Griffith et al. 2003). At the extreme, the organization is best served if the employee leaves after depositing his or her knowledge in the system, so that the organization can replace him or her with a new employee with different harvestable knowledge. The faster organizations can capture and store knowledge from various employees, the better they may perform, such that "in a Machiavellian world, organizations might develop systems where they quickly turn over employees after any unique knowledge has been stripped away" (Griffith et al. 2003, p. 280).

A third view suggests that moderate levels of turnover lead to the best organizational performance (Abelson and Baysinger 1984). Without turnover, the experience and knowledge of

organizational members become stagnant, obsolete, or overly insular (Dalton and Todor 1979; Shaw et al. 2005). When people leave, the organization likely hires new people, and moderate levels of turnover may create opportunities for organizations to obtain new skills and knowledge through the influx of new employees (Argote and Ingram 2000; Madsen et al. 2003). Although new members may have less experience than established members, their knowledge is typically less redundant with respect to the knowledge already possessed by the organization. New members thus might have a greater marginal impact on the knowledge held by the organization. Using this rationale, March (1991) finds that turnover relates in a curvilinear fashion to performance, such that moderate levels result in the highest levels of collaborative output.

Some research also suggests that moderate levels of turnover may benefit organizations that lack the time or resources to screen and select employees carefully (Siebert and Zubanov 2009). Turnover allows these organizations to be less discriminate in their hiring. Thus they can rely on moderate levels of turnover to retain the best employees and eliminate the worst (that is, after managers assess their on-the-job performance). Moderate turnover also improves performance if the detrimental impact of turnover is lower than the cost of eliminating it (Glebbeck and Bax 2004). What this means is that some turnover likely represents a natural state, and organizations thereby incur costs if they attempt to prevent or limit the amount of natural turnover (e.g., providing consistently challenging work, permitting job autonomy, etc.). The costs required to reduce turnover may exceed the negative effect of turnover, such that the optimal level of turnover is greater than zero.

Turnover and Performance in Online Communities

Given these differing views on the relationship between turnover and organizational performance, we next argue how membership turnover influences performance in social media

communities. The most common view of turnover in online communities, similar to that for organizations, is that it relates negatively to performance. The ability to attract and retain members frequently serves as a key metric for success in online communities (e.g., Arguello et al. 2006; Butler 2001; Lazar and Preece 2002; Ma and Agarwal 2007), because a stable group of participants can develop experience working together effectively, develop shared rules and norms, and agree on a common vision for the community (Lazar and Preece 2002; Ren et al. 2007). This shared experience might allow the community to work steadily toward a goal, whereas the loss of participants would mean that useful components of these shared norms and visions were no longer available to the community (Lazar and Preece 2002). Communities also tend to develop particular collaborative roles (e.g., content contributor, copy editors), and replacing these roles demands the time and energy of the remaining community members. They must find and train new members to perform these roles or else require existing members to perform these tasks, a reassignment that reduces the effort and energy available for fulfilling other roles they may have been performing. Finally, people often participate in online communities because they gain benefits from communicating and collaborating with others, so the departure of existing members may reduce the benefits of those left behind (Butler 2001).

Specific features of social media platforms may help mitigate some of these negative effects. Many social media platforms automatically store and retain contributions by participants, as a natural byproduct of the collaboration that occurs within the community (Kane and Fichman 2009). Thus, even when members leave social media communities, it does not necessarily follow that they take their knowledge with them. Instead, they leave behind a considerable part of the explicit knowledge that they have contributed, which then can be referenced, adapted, and used by the community. The automatic preservation of all information and communication also can

ensure tacit knowledge is owned by the community (Kane and Fichman 2009). When the entire collaborative history of a community is preserved, later members can use this history to discern effective norms, decision rules, and processes, even if the collaborators are not available to articulate those factors. In this sense, the platform may preserve and retain the experience gained through collaborating for the community that its individual members lack.

Furthermore, some turnover may be necessary to allow new members to join. Online communities are not technically limited to a finite size, but people tend not to join once the membership or communication levels are perceived to be too high (Butler et al. 2001; Kuk 2006). Groups that are isolated from outside perspectives can develop biases and insular thinking that leave them susceptible to overconfidence about the group's ability to collaborate effectively (Janis 1972; Schultze and Leidner 2002). Thus, some turnover might be necessary to create an influx of unique contributors with new ideas, skills, and information.

Although social media platforms may mitigate the negative effects of turnover, they are unlikely to eliminate them entirely. In particular, online communities must capture information contributed by participants and then organize it in a fashion that allows others to use it effectively (Alavi and Leidner 2001; Markus 2001; Stein and Zwass 1995). Markus (2001) describes such organizational processes as "culling, cleaning and polishing, structuring, formatting, or indexing documents against a classification theme" (p. 60). Uncritical collection and storage of all information actually makes it more difficult to identify the most important and relevant information (Hansen and Haas 2001). However, organization processes in online communities are often guided by norms and rules for effective collaboration (Butler et al. 2008), which are often developed by the community as they work together (Hinds and Bailey 2003). New participants may not be aware of or take advantage of the norms, rules, and history that the

platform provides, nor might the community ever develop sufficient norms in the presence of very high turnover.

We therefore expect that moderate levels of turnover are best for collaboration in online communities, such that membership turnover should have a curvilinear relationship with a community's performance. With too much turnover, prior knowledge generated by the community may be lost. With no ability to retain knowledge, the collaborative output of the community devolves into a random walk, only as valuable as the knowledge possessed by the most recent collaborators (Kane and Alavi 2007; March 1991). If too little turnover happens though, the knowledge created by the community can become stale and rigid (Garcia et al. 2003; Kane and Alavi 2007). In this case, the community may generate valuable knowledge and experience, but its value likely deteriorates over time. The community benefits from turnover to the extent that the influx of new knowledge exceeds the loss of existing knowledge held by departing members.

Online communities also typically have no mechanism to evaluate a member's potential before they join the community. Anyone can join an online community at any time, but the communities might impose processes to assimilate these new members slowly into the community, allowing them to become full active members only after an apprenticeship period (O'Mahony and Ferraro 2007; Preece and Schneiderman 2009). Members begin by observing the activity of the community, then start to contribute, and ultimately might end up as moderators and leaders of the community. Moderate levels of turnover may help the community identify and retain the best contributors, such that they move deeper into the community structure through an apprenticeship, much as traditional organizations rely on turnover to retain the best employees and remove the worst (Siebert and Zubanov 2009).

Multiple Stages of Collaboration in Online Communities

As we noted, collaboration in online communities may consist of different stages (Kane et al. 2009). First, the community must generate content. Tasks associated with this stage involve deciding which ideas should be included, refining those ideas to accommodate the multiple perspectives of participants, and integrating ideas with others developed in the community. Second, the community needs to maintain the relevance of the information it has generated. Neither collaboration nor membership turnover ceases simply because the community has generated high quality content, and the community must work actively to ensure that new collaboration does not destroy the information it has generated. Communities do not necessarily proceed through these stages linearly but rather may go through cycles of creation and maintenance as information improves incrementally (Kane 2011) or confront the need to recreate knowledge in the face of failed maintenance efforts (Kane et al. 2009).

We anticipate that membership turnover affects collaboration in similar ways during both stages of development. The stages represent a distinct shift in collaborative emphasis but not fundamentally different collaborative processes. Despite their differing goals, both stages demand a balance between incorporating new information and preserving the information the community already possesses (cf. Kane and Alavi 2007; March 1991). In the information-generating stage, the community may place a greater emphasis on new information, but it must retain some information previously generated, lest collaborative processes reflect only information possessed by the most recent collaborators. In the information retention stage, the community may place a greater emphasis on protecting extant information, but it still must integrate new information provided by new participants, lest the information become stale and obsolete in relation to a changing information environment. Thus, we expect that membership

turnover relates in a curvilinear way to performance during at the generating and retaining stages. We posit:

H1a: Membership turnover in an online community relates in a curvilinear fashion to knowledge creation, improving it up to an optimal point and impairing it thereafter

H1b: Membership turnover in an online community relates in a curvilinear fashion to knowledge retention, improving it up to an optimal point and impairing it thereafter.

RESEARCH METHOD AND SETTING

We test the impact of membership turnover on collaborative success by investigating the development of articles on Wikipedia. Wikipedia uses a wiki platform to host an open-source encyclopedia. Drawn from the Hawaiian word meaning “quick,” a wiki is simply a Web site that anyone can edit. Established in 2001, the English version of Wikipedia has developed, as of this writing, around 3.5 million separate articles. The scores of other languages that Wikipedia also hosts contain an additional 13 million articles. Anyone can edit any article on Wikipedia. When the user does so, the platform records the editor’s identity, the changes he or she has made to the article, a description of the change, and the time of the change. Other users can be automatically notified of any changes to a particular article, and they can undo any set of edits to previous versions. Representing a multi-year, collaborative development by public volunteers, Wikipedia is a robust setting for studying IT-enabled collaboration in online communities (Kane and Fichman 2009; Wagner and Majchrzak 2006).

The goal of collaboration on Wikipedia is to create information that is accurate, neutral, complete, and well written (<http://en.wikipedia.org/wiki/Wikipedia:Fa>). Wikipedia provides a rating system for articles to determine how well a particular article has achieved these goals. The editors of Wikipedia identify articles that exemplify these collaborative standards and award them the distinction of *featured article*. Featured articles are only small fraction of the total

articles (<0.1%) on Wikipedia, though anyone can nominate an article for promotion to this status. Once nominated, a forum reminiscent of an academic peer review vets candidate articles; reviewers highlight their strengths and weaknesses. Anyone can review nominated articles for a three-week period, after which the featured article director (a position similar to a senior editor at an academic journal) determines whether there is consensus for promotion. Unlike academic peer review, the process is open to anyone for comment.

Although featured article status might not be the goal to which all individual participants of the community aspire, it is the stated primary goal of Wikipedia. Therefore, we use featured article status not as a proxy for the creation of more knowledge that is objectively more valuable¹ but rather as a measure of whether the community has achieved the goal for which it was founded—much as a company might establish a community for a particular goal that may differ from the reasons that individual members participate.

When editors award featured article status, they do not guarantee it will persist. Anyone may demote an article to regular status at any time if he or she determines that it no longer meets the required standards. In this case, the assessment process is similar to the peer reviews that occur prior to promotion except the article is nominated for demotion rather than promotion. Because the community evaluates the collaborative standards of accuracy and completeness in relation to the current state of knowledge, neither collaboration nor membership turnover ceases once editors award featured article status. We qualitatively examined the promotion and demotion process of 100 articles from our sample. We identified each instance where an article was considered for demotion and then studied the transcripts associated with the deliberations to

¹ Research has validated the objective accuracy of featured article status. Kittur and Kraut (2008) demonstrate its objective validity using third-party reviewers, and Kane (2009) uses medical school students to substantiate the objective validity of ratings of Wikipedia's medical articles.

determine what factors led to demotion. We found that the articles in our sub-sample were nominated for demotion on average once per year following promotion, suggesting that the likelihood of demotion is a real and constant threat for Featured Articles, not a random or occasional process. The most common criteria for nominating an article for demotion was either that the information was somehow dated or irrelevant (i.e., no longer complete/accurate) or that newly added information detracted from the article's overall presentation of information (i.e., it was no longer well-written). This evidence suggests that effective collaboration needs to continue even after promotion, though the collaborative goals associated with maintaining the article might differ somewhat from collaborative goals associated with developing it.

Given this background on featured articles, we chose to use promotion to and demotion from featured article status as a surrogate for collaborative success. In the empirical testing, we first examined the influence of membership turnover on an article's likelihood of promotion (conditional on an eventual promotion), then test the influence of membership turnover on the likelihood that editors demote an article. Because featured article status is not permanent, articles may be promoted and demoted more than once, so we examined each phase separately.

Data

We built a 186-gigabyte data set of the full text of 3,720,826 revisions of 2,065 articles, from the inception of Wikipedia in 2001 until 2008. All selected articles had been featured articles at some point between 2001 and 2008. In our study, the online community was comprised of people who had contributed to a selected article during that time. Of the 2,065 articles that attained featured article status, editors eventually stripped 447 of the distinction. Because of the volume of these data and possible lag effects, we aggregated the individual revisions to 118,474 monthly observations of editing activity on the focal articles. Overall,

736,054 distinct authors revised the featured articles, and each author contributed an average of 5.06 revisions. In addition, automated programs called *bots* made many revisions automatically. In fact, there were 143 bots active during the study period for articles in our sample; these were responsible for 38,001 revisions, all of which we excluded from the analysis.

We measured membership turnover by recording the contributors previous experience collaborating on the article, the opposite condition of turnover. In communities with high turnover, contributors will have on average relatively low previous experience contributing to the content. In communities with low turnover, contributors will have on average relatively high previous experience. To construct the independent variable, *average experience*, we collected data from the logs preserved by the Wikipedia platform and averaged the experience of each contributor. For each revision, we built a measure of the experience of the author in the community by totaling the number of prior edits the author had made to the article. We then averaged experience by dividing this number by the total number of edits made to the focal article during the monthly observation. For example, an editor with no prior contributions to an article has an experience value of 0, whereas an editor who made 100 previous edits has an experience value of 100. We then summed the total volume of experience across all contributors and divided by the number of edits to yield an experience average. Higher average experience means that most contributions are made by members who have been in the community for a long time, which suggests low turnover conditions. Lower average experience means that most contributions are made by members relatively new to the community and thus could be characterized as high turnover.

Control Variables

Featured articles likely are 1) complete, 2) accurate, 3) well written, and 4) neutral. Therefore, we control for several alternative attributes of the article that may influence its characteristics by including additional variables in the model as predictors of promotion and demotion. First, the community may perceive longer articles as more complete because they have more content; therefore, we control for the total length of each article. *Article length* is the total number of characters of text, which ranges from 19 to 1,047,752, with an average of 25,152 characters. Similarly, the community may perceive articles with more complex organizational structures as more complete. Participants organize articles into sections, so we include a *section depth* variable to capture the organizational structure. It is the maximum depth of section levels, ranging from 1 to 6 with an average of 2.52 levels. (We also measured the total number of sections, which was highly correlated with the length of the article and therefore not included in the analysis.)

Second, the community may perceive articles that use more sources as more accurate. Articles must provide reliable sources (e.g., mainstream news media, research articles) for all original ideas contained in the article. We count the number of these *external references* per article, then divide by the article length to gain a measure of the intensity of references. Articles also may reference other articles on Wikipedia, typically if those articles provide a more extensive treatment of a related subject. An article with a greater number of links to other articles may appear more accurate, because it relies on other articles to address ancillary issues. We label these citations *internal references* and again count the number per article. Because the number of internal references necessarily correlates with article length, we divide the number of references by article length to develop a measure of the intensity of internal citations.

Third, stylistic differences may cause articles to appear better written. To control for possible stylistic differences, we consider reading complexity and the use of multimedia images. For each revision, we measure the reading *complexity* of the article using the automated readability index (ARI; Smith and Senter 1962)². In addition, multimedia images may affect the perceived writing quality, so we include a measure of *multimedia intensity*, for which we count the number of multimedia files associated with each revision. Because this measure correlates with article length, we again divide the count by the article length.

Fourth, the number of edits could be associated with the article’s completeness and perceived neutrality. Highly controversial articles likely provoke considerable dispute, especially about sensitive terms, and more edits may mean that the community has spent considerable time in the editing process. Appendix A provides a descriptive overview of these data and the correlations between our focal variables.

Data Analysis

Our sample includes all articles in Wikipedia that have achieved featured article status at some point during their existence. We evaluate the influence of our focal independent variable, community experience, on the likelihood of two separate events. First, we examine the promotion of articles to featured article status. Second, we examine the likelihood of demotion from featured article status. Our sampling strategy enables us to compare the same set of articles across both stages but does not allow us to estimate the likelihood of promotion of *any* Wikipedia article—we can predict only those that eventually are promoted. The interpretation of

² The formula is $ARI = \left(\left(4.71 \times \frac{\text{letters}}{\text{words}} \right) + \left(\frac{0.5 \times \text{words}}{\text{sentences}} \right) - 21.43 \right)$. The empirically derived index estimates

the U.S. School grade required to understand a text. As a robustness check, we also compared the models using the Coleman-Liau index and found no substantive differences.

the coefficients in both the control and independent variables will be opposite in our models, because the likelihood of promotion is a desirable outcome, whereas the likelihood of demotion is an undesirable one.

We evaluate the likelihood of the two events using semi-parametric proportional hazard modeling. Proportional hazard models are useful for assessing the effect of a measure on the likelihood of an event occurring. They require no assumption of the functional form of the underlying likelihood of the event; instead, they assume that variables have a proportional effect on the unspecified underlying likelihood. Therefore, we can assess how our focal variable, community experience, changes the likelihood that the article will experience a promotion or demotion event. Mathematically, we describe the rate (λ) that an article will experience an event (promotion or demotion) at time t as

$$\lambda(t) = \lim_{\Delta t \rightarrow 0} \frac{\Pr(t \leq T \leq t + \Delta t | T \geq t)}{\Delta t},$$

assuming the article has not experienced the event prior to time T . The general model $\lambda(t_i) = \lambda_0(t_i)e^{-x_i\beta}$ serves to estimate the coefficients of impact (β) on a baseline hazard rate (λ_0), resulting in a hazard rate (λ) at time t_i for *featured article i* (Cox 1972). In Appendix B, we offer further details about the data analysis methods and robustness checks.

RESULTS

Table 3 contains the results of our empirical analysis. For the analysis, we standardized all continuous variables by subtracting their mean and dividing by the standard deviation, which unified the presentation of the results and facilitated comparisons.

Models 1 and 2 test our variables in relation to the rate of promotion. Model 1 includes only the control variables, and Model 2 adds the linear and squared average experience. Both the

linear ($\beta = 0.974, p < 0.01$) and squared ($\beta = -1.164, p < 0.01$) coefficients are significant and in the expected direction. Membership turnover, measured by the average amount of contributors' previous experience, has a curvilinear relationship with the likelihood of article promotion. Moderate levels of membership turnover, captured by a mid-range level of average experience, lead to the greatest chance of article promotion. Models 3 (control variables) and 4 (with linear and squared average experience) then demonstrate the results in relation to the likelihood of demotion. Again, both the linear ($\beta = -0.869, p < 0.01$) and squared ($\beta = 0.323, p < 0.01$) coefficients are significant and in the expected direction. Membership turnover thus has a curvilinear relationship with article demotion, such that moderate levels lead to the lowest probability of demotion.

Table 1: Proportional Hazard Analysis of the Effect of Membership Turnover

| | Promotion | | Demotion | |
|--------------------------------|-----------|----------|----------|----------|
| | Model 1 | Model 2 | Model 3 | Model 4 |
| ‡ Article length | | | -0.169 | -0.149 |
| ‡ Section depth | | | -0.001 | -0.011 |
| ‡ External references | | | -1.566** | -1.484** |
| ‡ Internal references | | | -0.058 | -0.058 |
| ‡ Reading complexity | | | -1.193 | -1.224 |
| ‡ Number of edits | | | 0.051 | 0.086 |
| ‡ Experience/Edit | | | -0.327** | -0.221** |
| Article length | 0.108** | 0.107** | -0.024 | -0.023 |
| Section depth | 0.624** | 0.487** | 0.036 | 0.038 |
| External references | 0.397** | 0.244** | -0.168** | -0.162** |
| Internal references | 0.042 | 0.065** | 0.420** | 0.408** |
| Reading complexity | -0.010 | -0.004 | -0.073 | -0.078 |
| Media references | 0.027* | 0.040** | -0.013 | -0.012 |
| Number of edits | 0.059 | -0.002 | 0.144** | 0.184** |
| Average experience | | 0.974** | | -0.869** |
| (Experience/Edit) ² | | -1.624** | | 0.323** |
| Observations | 64,554 | 64,554 | 47,364 | 47,364 |
| Log likelihood | -12541 | -12197 | -2708 | -2694 |
| χ^2 | 1529.33 | 1623.65 | 243.83 | 250.56 |
| Adjusted pseudo R ² | 8.8 | 11.3 | 7.3 | 7.8 |

| | | |
|--------------------------------|----------------------|---------------------|
| Effect size (f2) of experience | 0.03 | 0.01 |
| Pseudo F values | 1821.29 ⁺ | 239.21 ⁺ |

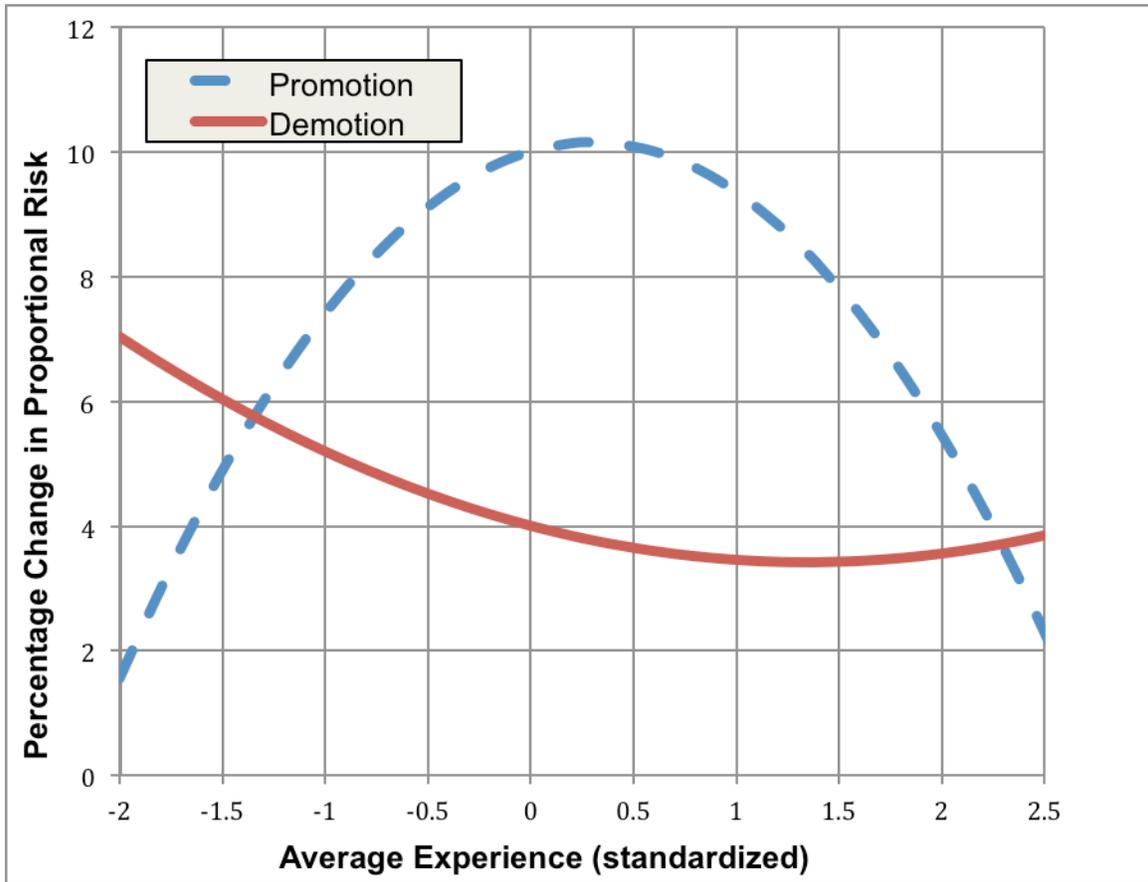
‡ Value at promotion

*Significant at < 0.05. **Significant at < 0.01.

⁺ Effect size pseudo-F values are significant at an alpha level of 0.001.

Because our promotion and demotion models value the event in opposite directions (e.g., promotion is good, demotion is bad), the coefficients of the squared parameters also should move in opposite directions. Visualizing these results may aid in their interpretation, so in Figure 1, we depict the curvilinear relationship between experience and collaborative success. For the promotion stage, our dependent variable is the likelihood that an article will be promoted to featured article status, which is a positive outcome. As Figure 1 shows, average experience increases the likelihood of promotion up to a certain point but then decreases the likelihood. Thus, the shape of the promotion curve, such that mid-range levels of turnover enjoy the greatest chance of promotion, supports H1a. For the demotion stage, our dependent variable is the likelihood that an article will be demoted from featured article status, a negative outcome. The graphs in Figure 1 show that average experience decreases the likelihood of demotion up to a certain point, after which it begins increasing that likelihood. Thus, the shape of the demotion curve, according to which mid-range levels of experience enjoy the lowest chance of demotion, supports H1b.

Figure 1: Relationship Between Experience and Collaborative Success



Our large sample size also prompts us to discuss the effect size of our results (Burton-Jones and Straub 2006; Chin et al. 2003). First, Model 2 has a pseudo- R^2 value of 11.3 (see the Appendix), with an effect size (0.03) that falls between small and medium (Cohen 1988); based on the pseudo-F value, it represents a significantly improved model ($\alpha = 0.001$). The comparison of the pseudo- R^2 of Models 1 and 2 shows that including our turnover variables explains nearly 30% more of the variance in likelihood than the baseline model. Furthermore, the coefficient of the standardized variables reveals that the experience variables have a far greater effect than any of the other variables in our model. Thus, membership turnover is clearly an

important variable with respect to explaining the knowledge creation stage of collaboration in social media communities.

Second, in the demotion models, the increase in variance explained is smaller (~7%), and the effect size (0.01) is in the small range (Cohen 1988); however, based on the pseudo-F value, the inclusion of turnover still significantly improves the model ($\alpha = 0.001$). It is perhaps not surprising that the change in the pseudo- R^2 that we observe with the demotion model is smaller than that for the promotion model when we add the turnover variables because these statistics are sensitive to the large number of non-failure events (Harrell 2001). That is, most articles are not demoted. Nevertheless, the standardized coefficients show that these variables exhibit among the largest effect sizes of any variables in the model, so the level of turnover is important for the knowledge retention phase of collaboration in social media communities.

When considering effect size, we must also consider the real-world impact of certain effects. For example, even a 0.01% change in the U.S. gross domestic product still represents a significant real-world change, on the magnitude of billions of dollars. Wikipedia is one of the most heavily trafficked sites on the Internet, and changes to featured articles' status may have considerable effects on how often the content gets viewed and the degree to which people rely on this information. To assess the effect of featured status on viewership, we examined a sample of the 14,088 articles in the Wikipedia Medicine project from December 2007 until March 2009. Featured articles in this sample had more than seven times the viewership than non-featured articles; on average featured articles were viewed 39,918 times per month whereas non-featured articles were viewed 5,691 times per month. (The difference is statistically significant with a t value of 16.06 for the differences in monthly average viewership.) Thus, even relatively modest changes in the likelihood of promotion or demotion have significant real-world impacts on the

adoption and diffusion of information created by social media communities. Certainly companies seeking to create information in peer-production communities need insight into how to improve or preserve the value of information they create.

It also may be worth examining the shapes and inflection points of the curves in Figure 1. Although the curve associated with the promotion model is far more pronounced, the inflection point is near 0. Thus, the ideal amount of turnover is approximately equal to the mean levels possessed by the articles in our sample. In contrast, the curve associated with the demotion model is less pronounced, but the inflection point is 1–1.5 standard deviations above the norm. That is, the typical article experiences more turnover (i.e., lower experience/edit) than is ideal during the retention stage. Members of the community appear to focus on creating the content but show less interest in preserving that content. The community might generally have determined an appropriate mix for optimizing the content generation phase; for the relatively newer task of knowledge retention though, the typical community is experiencing too much turnover.

Differences in Knowledge Creation and Retention

Because little is known about collaboration on the Wikipedia platform, we examine some of the results related to our control variables in Table 3. These results provide some further empirical evidence that knowledge creation and retention are separate stages of collaborative development in social media communities. Although many of the variables associated with article content—length, section depth, multimedia references, external references, and internal references—are positively associated with the likelihood of promotion, very few of them significantly affect the likelihood of demotion. Better referenced articles developed by communities with more experience may withstand potential demotion somewhat better, but it

appears that the factors associated with collaboration after promotion actually have a greater effect on the likelihood of demotion.

Furthermore, the factors associated with desired outcomes at both stages differ. Only the effects of external references are consistent for both promotion and demotion phases; content that offers robust and timely references is generally considered superior. Although internal references are significant at both stages, they relate in contrasting ways to knowledge creation and retention. Perhaps in the knowledge creation stage, collaborators use internal links to legitimize content by associating it with other content, such that the contributors seek legitimacy for their content by connecting to other quality content. In the knowledge retention stage though, the opposite effect might occur. Since anyone can contribute to the content, outside contributors might seek use the legitimacy of the focal content to increase the legitimacy of or traffic to their own content by creating links from the focal article. These links in the knowledge retention stage are thus created for purposes other than improving the quality of the focal content.

Finally, though the number of edits is not associated with promotion, it is associated with demotion; content that is edited more heavily after promotion is more likely to be demoted. We do not interpret this result to mean that further collaboration should not occur after knowledge has been created, because new, well-referenced material is clearly important for keeping the content up-to-date. However, this finding lends some credence to the saying, “if it ain’t broke, don’t fix it.” The collaborative challenge in the knowledge retention phase may be mostly associated with protecting the integrity of content, rather than significant ongoing content generation.

DISCUSSION AND CONCLUSION

This article investigates the collaboration associated with 2,065 Wikipedia articles leading up to and following their promotion to featured article status to assess how membership turnover may be associated with success during both the knowledge creation and knowledge retention stages of collaboration. We find that membership turnover has a curvilinear effect on success in both stages.

Theoretical Implications

These findings have several important implications for collaboration in online communities. Most previous work on this topic has assumed that membership retention is a positive condition for online communities (e.g., Arguello et al. 2006; Butler 2001; Lazar and Preece 2002; Ma and Agarwal 2007), which may have been true for previous generations of online communities but does not appear accurate in reference to social media communities. Our results indicate that moderate levels of membership turnover are desirable in social media communities, because such levels offer new information and abilities to the community, without compromising its ability to retain the content it has generated. Further research into social media communities therefore should not assume that membership turnover is necessarily an undesirable characteristic for collaboration. We do, however, find evidence that communities on average experience more turnover during the knowledge retention phase than would be optimal for effective collaboration.

We also find empirical evidence to support the understanding that knowledge creation and knowledge retention are distinct phases in community-based peer production. The collaboration that occurs during the knowledge creation phase has little bearing on the effectiveness of knowledge retention; and the factors associated with effective knowledge

retention are also different from those associated with effective knowledge creation. This finding suggests that future research into community-based peer production should consider the state of the production process and recognize that the characteristics of and objectives for collaboration may differ based on the stage of production.

Managerial Implications

This investigation also has implications for the managers leading or managing community-based peer production environments. These managers might seek intentionally to cultivate a core group of members who participate over the long term, such as by offering incentives to a small number of participants (e.g., employees, customers) who agree to remain active in the community. However, they also should encourage the community to remain open to outsiders, who can join and leave at will. Such outsiders do not necessarily need to remain active in the community; rather, the long-term members and manager should find ways to organize and preserve their contributions, even if they leave. It also may be necessary for members of the core group to leave eventually, which allows new members to assume leadership roles and introduce new resources to meet the changing collaborative needs of the community. Managers should also recognize that the collaborative challenges for communities focused on information generation may entail different elements than those of communities focused on information retention, modifying their leadership style and goals in relation to those appropriate the stage of community's production.

Limitations

This paper contains several limitations that influence the potential generalization of its findings. First, we conducted this research entirely within the Wikipedia environment; additional

research will be necessary before the results can be generalized to other social media communities. Specifically, other social media platforms store and present the collaborative activities of their members in different ways, which may affect the relationship between turnover and performance. Other social environments, such as corporate social media communities or electronic networks of practice, may create different social conditions that lead to different levels or effects of turnover. In the Wikipedia context, many roles (e.g., copy editing, subject matter expertise, community understanding) are important at different times, whereas the importance of membership turnover likely varies in environments with more limited or expanded roles and more temporal or enduring interactions.

Second, we focus on a particular set of high-quality articles on Wikipedia. The vast majority of articles never reach featured article status. Therefore, further research should explore whether these findings hold in less well-developed collaborative environments.

Third, to compare collaboration across the creation and retention stages, our promotion models feature only articles that eventually get promoted to featured article status. Researchers also might examine a broader sample of articles (e.g., all articles nominated for featured article status) in the creation stage to test the robustness of our findings.

Conclusion

Despite these limitations, this paper makes important contributions. In particular, it provides empirical evidence that moderate levels of membership turnover positively affect collaborative success. Some membership stability is necessary to retain the information and knowledge generated by the community, but turnover also is desirable to introduce new information to the community. Moreover, knowledge retention in online communities is fundamentally different from knowledge creation. Thus, this study offers several insights that

extend our understanding of the under-investigated phenomenon of knowledge retention in online communities.

ACKNOWLEDGEMENTS

The authors wish to acknowledge funding for this research by the National Science Foundation (CAREER 0953285). The authors also wish to recognize the Research Committee of the Carroll School of Management for their generous support of the 1st Boston College Social Media Workshop, and the participants in that workshop for helpful comments on this paper in its earliest stages.

REFERENCES

- Abelson, M.A., and Baysinger, B.D. "Optimal and dysfunctional turnover: Toward an organizational level model.," *Academy of Management Review* (9) 1984, pp 331-341.
- Alavi, M., and Leidner, D.E. "Review: Knowledge management and knowledge management systems: Conceptual foundations and research issues," *MIS Quarterly* (25:1), Mar 2001, pp 107-136.
- Argote, L., and Epple, D. "Learning curves in manufacturing.," *Science* (247) 1990, pp 920-924.
- Argote, L., and Ingram, P. "Knowledge transfer: A basis for competitive . advantage in firms," *Organizational Behavior Human Decision Processes* (82) 2000, pp 150-169.
- Arguello, J., Butler, B.S., Joyce, E., Kraut, R., Ling, K.S., Rose, C., and Wang, X. "Talk to me: foundations for successful individual-group interactions in online communities," CHI '06 Proceedings of the SIGCHI conference on Human Factors in computing systems, 2006.
- Becker, G.S. "Investment in human capital—a theoretical analysis," *Journal of Political Economy* (70:5) 1962, pp 9-49.
- Burton-Jones, A., and Straub, D.W. "Reconceptualizing system usage: An approach and empirical test," *Information Systems Research* (17:3), Sep 2006, pp 228-246.
- Butler, B.S. "Membership size, communication activity, and sustainability: A resource-based model of online social structures," *Information Systems Research* (12:4), Dec 2001, pp 346-362.
- Butler, B.S., Joyce, E., and Pike, J. "Don't look now, but we've created a bureaucracy: the nature and roles of policies and rules in Wikipedia," 2008. Proceeding of the twentieth-sixth annual SIGCHI conference on human factors in computing systems.

- Butler, J., Morrice, D.J., and Mullarkey, P.W. "A multiple attribute utility theory approach to ranking and selection," *Management Science* (47:6), Jun 2001, pp 800-816.
- Carley, K. "Organizational learning and personnel turnover," *Organization Science* (3:1), Feb 1992, pp 20-46.
- Chin, W.W., Marcolin, B.L., and Newsted, P.R. "A partial least squares latent variable modeling approach for measuring interaction effects: Results from a Monte Carlo simulation study and an electronic-mail emotion/adoption study," *Information Systems Research* (14:2), Jun 2003, pp 189-217.
- Cohen, J. *Statistical Power Analysis for the Behavioral Sciences*, (2nd ed.) Lawrence Erlbaum, Hillsdale, NJ, 1988.
- Cox, D.R. "Regression models and life tables," *Journal of the Royal Statistical Society, Series B* (34) 1972, pp 187-202.
- Dalton, D.R., and Todor, W. "Turnover turned over: An expanded and positive perspective," *Academy of Management Review* (4:2) 1979, pp 225-235.
- Darmon, R.Y. "Identifying sources of turnover costs—a segmental approach," *Journal of Marketing* (54:2), Apr 1990, pp 46-56.
- Dess, G.G., and Shaw, J.D. "Voluntary turnover, social capital, and organizational performance," *Academy of Management Review* (26:3), Jul 2001, pp 446-456.
- DiGangi, P., Wasko, M., and Hooker, R. "Getting customers' ideas to work for you: Learning from Dell how to succeed with online user innovation communities," *MIS Quarterly Executive* (9:4) 2010.

- Faraj, S., Jarvenpaa, S., and Majchrzak, A. "Knowledge collaboration in online communities," *Organization Science*, forthcoming, published online before print Feb 23, 2011 , DOI: doi:10.1287/orsc.1100.0614.
- Gallaughar, J., and Ransbotham, S. "Social media and customer dialog management at Starbucks," *MIS Quarterly Executive* (9:4) 2010.
- Garcia, R., Calantone, R., and Levine, R. "The role of knowledge in resource allocation to exploration versus exploitation in technologically oriented organizations," *Decision Sciences* (34:2), Spr 2003, pp 323-349.
- Glebbeeck, A.C., and Bax, E.H. "Is high employee turnover really harmful? An empirical test using company records," *Academy of Management Journal* (47:2), Apr 2004, pp 277-286.
- Griffith, T.L., Sawyer, J.E., and Neale, M.A. "Virtualness and knowledge in teams: Managing the love triangle of organizations, individuals, and information technology," *MIS Quarterly* (27:2), Jun 2003, pp 265-287.
- Hansen, M.T., and Haas, M.R. "Competing for attention in knowledge markets: Electronic document dissemination in a management consulting company," *Administrative Science Quarterly* (46:1), Mar 2001, pp 1-28.
- Harrell, F. *Regression Modeling Strategies* Springer-Verlag, New York, 2001.
- Hinds, P.J., and Bailey, D.E. "Out of sight, out of sync: Understanding conflict in distributed teams," *Organization Science* (14:6), Nov-Dec 2003, pp 615-632.
- Hom, P., and Griffeth, R. *Employee Turnover* South-Western Publishing, Cincinnati, OH., 1995.

- Huselid, M.A. "The impact of human resource management practices on turnover, productivity, and corporate financial performance," *Academy of Management Journal* (38:3), Jun 1995, pp 635-672.
- Janis, I. *Victims of Groupthink* Houghton Mifflin Company, Boston, 1972.
- Kane, G.C. "It's a network, not an encyclopedia: A social network perspective on Wikipedia collaboration," Academy of Management Annual Meeting Proceedings, Chicago, IL, 2009.
- Kane, G.C. "A multi-method study of information quality in Wiki collaboration," *ACM Transactions on Management Information Systems* (1:2) 2011.
- Kane, G.C., and Alavi, M. "Information technology and organizational learning: An investigation of exploration and exploitation processes," *Organization Science* (18:5) 2007, pp 796-812.
- Kane, G.C., and Fichman, R.G. "The shoemaker's children: Using Wikis for IS research, teaching, and publication," *MIS Quarterly* (33:1) 2009, pp 1-22.
- Kane, G.C., Majchrzak, A., Johnson, J., and Chen, G.L. "A lifecycle model of perspective making and perspective taking in fluid online collectives," in: *International Conference on Information Systems*, Phoenix, AZ, 2009.
- Kittur, A., Chi, E., Pendleton, B.A., Suh, B., and Mytkowicz, T. "Who writes Wikipedia? How centralization scales (and doesn't) in social collaborative systems.," Working Paper, 2007.
- Krackhardt, D., and Porter, L.W. "When friends leave: A structural analysis of the relationship between turnover and stayers' attitudes," *Administrative Science Quarterly*, 1985, p. 242.

- Kuk, G. "Strategic interaction and knowledge sharing in the KDE developer mailing list," *Management Science* (52:7), Jul 2006, pp 1031-1042.
- Lakhani, K.R., and Kanji, Z. "Threadless: The business of community," Harvard Business School Case, Cambridge, MA, 2008.
- Lazar, J., and Preece, J. "Online communities: Usability, sociability and users' requirements," in: *Cognition in the Digital World.*, H.v. Oostendorp (ed.), Lawrence Erlbaum Associates Inc. , Mahwah, NJ, 2002, pp. 127-151.
- Leana, C.R., and Van Buren, H.J. "Organizational social capital and employment practices," *Academy of Management Review* (24:3), Jul 1999, pp 538-555.
- Li, C., and Bernoff, J. *Groundswell: Winning in a World Transformed by Social Technologies* Harvard Business Press, Boston, MA, 2008, 286 p.
- Ma, M., and Agarwal, R. "Through a glass darkly: Information technology design, identity verification, and knowledge contribution in online communities," *Information Systems Research* (18:1), Mar 2007, pp 42-67.
- Madsen, T.L., Mosakowski, E., and Zaheer, S. "Knowledge retention and personnel mobility: The nondisruptive effects of inflows of experience," *Organization Science* (14:2), Mar-Apr 2003, pp 173-191.
- March, J.G. "Exploration and exploitation in organizational learning," *Organization Science* (2:1), Feb 1991, pp 71-87.
- Markus, M.L. "Toward a theory of knowledge reuse: Types of knowledge reuse situations and factors in reuse success," *Journal of Management Information Systems* (18:1), Sum 2001, pp 57-93.

- Nelson, R., and Winter, S. *An Evolutionary Theory of Economic Change*. Harvard University Press, Cambridge MA, 1982.
- O'Mahony, S., and Ferraro, F. "The emergence of governance in an open source community," *Academy of Management Journal* (50:5), Oct 2007, pp 1079-1106.
- Oh, W., and Jeon, S. "Membership herding and network stability in the open source community: The Ising perspective," *Management Science* (53:7), Jul 2007, pp 1086-1101.
- Preece, J., and Schneiderman, B. "The reader-to-leader framework: Motivating technology-mediated social participation," *AIS Transactions on Human-Computer Interactions* (1) 2009, pp 13-32.
- Ren, Y., Kraut, R., and Kiseler, S. "Applying common identity and bond theory to design of online communities.," *Organization Studies* (38) 2007, p 377+.
- Sarner, A. "The business impact of social computing on marketing and 'generation virtual'." Gartner Symposium/ITxpo October 12-16, 2008.
- Schemper, M., and Stare, J. "Explained variation in survival analysis," *Statistics in Medicine* (15) 1996, pp 1999-2012.
- Schultze, U., and Leidner, D.E. "Studying knowledge management in information systems research: Discourses and theoretical assumptions," *MIS Quarterly* (26:3), Sep 2002, pp 213-242.
- Shaw, J.D., Gupta, N., and Delery, J.E. "Alternative conceptualizations of the relationship between voluntary turnover and organizational performance," *Academy of Management Journal* (48:1), Feb 2005, pp 50-68.

- Siebert, W.S., and Zubanov, N. "Searching for the optimal level of employee turnover: A study of a large UK retail organization," *Academy of Management Journal* (52:2), Apr 2009, pp 294-313.
- Smith, E. A. and Senter, R. J. "Automated readability index." Aerospace Medical Research Laboratories, Wright-Patterson Airforce Base, 1962. pp 1-14.
- Staw, B.M. "The consequences of turnover," *Journal of Occupational Behaviour* (1:4) 1980, pp 253-273.
- Stein, E.W., and Zwass, V. "Actualizing organizational memory with information systems," *Information Systems Research* (6:2), Jun 1995, pp 85-117.
- Ton, Z., and Huckman, R.S. "Managing the impact of employee turnover on performance: The role of process conformance," *Organization Science* (19:1), Jan-Feb 2008, pp 56-68.
- Wagner, C., and Majchrzak, A. "Enabling customer-centricity using wikis and the wiki way," *Journal of Management Information Systems* (23:3), Win 2006, pp 17-43.

APPENDIX A

Descriptive Statistics

| Variable | Minimum | Maximum | Mean | Std. Dev |
|---|---------|---------|--------|----------|
| 1. Length ($\times 10^6$) | 0 | 1.048 | 0.029 | 0.036 |
| 2. Section depth | 1 | 6 | 2.667 | 0.826 |
| 3. Internal references ($\times 10^3$) / length | 0 | 186.549 | 8.322 | 4.296 |
| 4. External references ($\times 10^3$) / length | 0 | 1616 | 27.081 | 45.507 |
| 5. Complexity | 0 | 335.971 | 0.035 | 1.506 |
| 6. Multimedia / length ($\times 10^3$) | 0 | 2.000 | 0.002 | 0.030 |
| 7. Edits | 0 | 15.591 | 0.558 | 0.949 |
| 8. Average experience | 0 | 0.942 | 0.019 | 0.041 |

Variable Correlations

| Variable | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
|---------------------------------|--------|--------|--------|--------|--------|--------|-------|
| 1. Length | 1.000 | | | | | | |
| 2. Section depth | 0.414 | 1.000 | | | | | |
| 3. Internal references / length | -0.291 | -0.338 | 1.000 | | | | |
| 4. External references / length | 0.492 | 0.367 | -0.285 | 1.000 | | | |
| 5. Complexity | 0.064 | 0.006 | -0.009 | 0.003 | 1.000 | | |
| 6. Multimedia / length | -0.027 | -0.031 | 0.214 | -0.024 | -0.001 | 1.000 | |
| 7. Edits | 0.439 | 0.384 | -0.164 | 0.470 | 0.008 | -0.241 | 1.000 |
| 8. Average experience | 0.309 | 0.270 | -0.201 | 0.501 | 0.001 | -0.207 | 0.292 |

APPENDIX B

The rich data set provides detailed observations of each focal article. For each article, we build a series of monthly observations of each variable, but our results are robust to expanding the window to both two and three months. Because we have multiple observations for each article, we can focus our analysis on changing editorial activity rather than unchanging attributes, such as topic. We include monthly spells of each article, either from the time of creation to promotion to *featured article* status or from the time of promotion to demotion or the end of our study period (December 2008). This censoring does not affect the promotion models, but for the demotion models, we incorporate censoring caused by the end of the study period.

Within the data set, we have multiple observations of each article—one for each monthly spell. Therefore, we allow variance to cluster by article to obtain robust variance estimates that control for possible within-article correlation. To control for editing prior to promotion in the demotion models, we include the values of all control variables at the time of article promotion. Our results are robust to including or excluding these controls. In addition, we tested shared frailty models, which represent a type of model misspecification correction that take their name from their use in tests of individual survival. An individual may have unobserved frailty that contributes to the likelihood of a failure. Because we have repeated monthly observations of each article, we can estimate an additional parameter to incorporate unobserved article-level heterogeneity. However, tests using shared frailty models reveal no variance in frailty θ and are significantly different from zero. Therefore, we retain and report on simpler, unshared frailty models.

Several metrics have been proposed for quantifying the explained variance in hazard models, though research on these models is ongoing (Schemper and Stare 1996). For each

model, we report a unit-less measure of predictive ability, analogous to an ordinary least squares adjusted R^2 . We calculate it as $pseudo - R^2 = ((L^0 - L) - 2p) / L^0$, where p is the number of coefficients estimated, L is the log likelihood of the focal model, and L^0 is the log likelihood from a baseline model with no coefficients estimated (Harrell 2001). Although imperfect, this measure permits some quantification of the variance explained by the model.

Using the change in pseudo- R^2 , we can estimate an effect size (f^2) by dividing the change in pseudo- R^2 by the unexplained variance ($1 - pseudo-R^2$) in the full model (Cohen 1988). We estimate the significance of the effect sizes by multiplying the effect size (f^2) by $(n - k - 1)$, where n is the sample size, and k is the number of variables in the full model. This calculation yields a pseudo-F value that we test using an alpha of 0.001.

Post-estimation diagnostics using Schoenfeld residuals indicate that our focal variables do not violate proportional hazard assumptions. However, the number of edits offers evidence of variance over time. In additional robustness checks, we find that our results are robust to models that permit the number of edits to vary linearly with time.

ONLINE SUPPLEMENT

We chose to measure turnover by the sum of the prior editing experiences. This measure captures the amount of experience from which each revision to an article was able to benefit. However, each revision to an article strictly increases the experience (since each revision cannot decrease the cumulative experience). Therefore, we need to isolate the cumulative experience from the number of edits; we do this by normalizing experience by the number of edits. Conceptually, a low value of this number indicates editing by relatively new (inexperienced) editors, while a high value indicates editing by relatively tenured (experienced) editors. We considered three alternative measures of experience.

- We could have used each term (sum of experience and number of edits) independently. However, using the two terms separately would not isolate the number of edits from the experience. Because experience comes from editing, the terms would be inherently correlated. In our context, both the numerator and denominator are important. It is their ratio which indicates the average experience of all the editing; neither the numerator nor denominator alone measures that.
- We could have used the standard deviation of user experience. However, it would capture the variation in experience between editors. We would therefore be unable to distinguish between a set of high experience editors (with variation in experience) and a set of low experience editors (with the same variation in experience). This would not capture turnover.
- We considered using Blau or Herfindal measures of diversity. However, they only capture the diversity in contributors; for example, they would not be able to distinguish between two groups of contributors that are similar with no experience and contributors that are similar with abundant experience.