It's a Network, Not an Encyclopedia: A Social Network Perspective on Wikipedia Collaboration

Authors

Gerald C Kane, Boston College, gerald.kane@bc.edu

It's a Network, Not an Encyclopedia: A Social Network Perspective on Wikipedia Collaboration

Abstract: This paper studies the collaboration process on Wikipedia to determine whether particular collaborative processes are associated with article quality. Employing a sample of 300 articles on medical and health-related topics, this paper examines the impact of the article's position within the two-mode affiliation network of articles and editors on article quality. This paper finds that the position of the article within both the local and global networks of articles and editors is significantly related to article quality. These results suggest that editors transfer information and knowledge from one collaborative environment to another on the same platform. Somewhat surprisingly, the direct collaborative processes have little relationship to article quality. Taken together, these findings suggest that the social aspects of Wikipedia are indispensable for understanding its collaborative processes.

Keywords: wiki, social networks, collaboraiton

Introduction

Community-based peer production is becoming an increasingly viable and popular way of creating knowledge and value for many organizations (Benkler, 2006; Lee & Cole, 2003; Raymond, 1999; von Hippel & von Krogh, 2003; Wagner & Majchrzak, 2006). Whether within, across, or even entirely independent from organizational boundaries, a new generation of Internet-based collaborative tools such as blogs, wikis, and electronic social networks are allowing people to interact with one another in new and potentially more powerful ways. These tools are used to allow employees to collaborate more effectively with one another, to allow employees to collaborate with dedicated customers, and even to create customer communities that interact primarily with one another regarding a company or product (Howe, 2008; Li & Bernoff, 2008; Tapscott & Williams, 2006). Although the technological features of these tools are critically important for enabling these collaborative practices, the interpersonal aspects of collaboration can be equally if not more important. This new generation of Internet-based collaborative tools can be deployed to address a particular organizational problem, but unless individuals actually use these tools to create effective collaborative practices as a result, these tools will create little valuable knowledge. Understanding the collaborative features that are associated with successful use of these tools can help organizations better leverage these tools to support collaboration.

Since it retains vast amounts of data on actual collaborative projects, one valuable environment for exploring these collaborative processes is Wikipedia (Kane & Fichman, forthcoming). Popularly known as the "encyclopedia that anyone can edit," Wikipedia is becoming an increasingly important example of community-based knowledge creation. As of

this writing, Wikipedia is one of the 10 most heavily trafficked Website on the Internet, contains over 10M articles in 40 languages, and its content has been repeatedly shown to be of surprisingly high quality (Devgan, Powe, Blakey, & Makary, 2007; Giles, 2005). This paper addresses seeks to identify whether certain collaborative features are associated with article quality on Wikipedia. Although care must be taken in generalizing collaboration effectiveness beyond the narrow confines of the Wikipedia environment, understanding collaborative success in this environment can provide managers insight into how to create similar types of collaborative environments in their own setting.

A growing stream of research has sought to investigate features of successful collaboration on Wikipedia (Kittur & Kraut, 2008; Wagner et al., 2006; Wilkinson & Huberman, 2007; Zhang & Zhu, 2007). This paper represents an important departure from this earlier work, as it adopts fundamentally different assumptions about the collaborative environment. Most previous research into collaboration on Wikipedia has approached it from an article-level perspective, treating the collaborative efforts associated with individual articles as independent from one another. In contrast, this paper does not assume that articles are independent. It fact, it explicitly explores whether and how the connections *between* articles influence article quality. Wikipedia editors often collaborate with a variety of different editors to develop multiple articles. As editors work on multiple articles on the Wikipedia platform, they can serve as conduits to transfer information and knowledge from one collaborative effort (i.e. article) to another. As such, the quality of a particular article may depend not only on the collaborative activity related to the development of that article but also on how that article is situated via the other articles on Wikipedia. This interpretation would suggest that the encyclopedia metaphor

that has come to define Wikipedia is somewhat misleading, and that Wikipedia is perhaps better regarded as an interconnected network of knowledge than as a collection of independent articles.

Using data from the collaborative activity of 300 articles from the Medicine Wikiproject, a sub-community on Wikipedia dedicated to the development of medical information, this paper adopts a social network approach to develop and test two hypotheses regarding whether two different measures of an article's centrality in the network of articles and editors – degree centrality and eigenvector centrality – will be positively related to article quality. This paper finds that both centrality measures are strongly associated with article quality. Combined with the finding that very little of the direct collaborative activity related to a given article is associated with article quality, this paper suggests that collaboration that occurs *between* collaborative environments is as important for understanding article quality as the collaboration that occurs within any given environment.

Social Network Analysis

One productive approach to studying collaborative environments has been through the application of social network analysis (Borgatti & Cross, 2003; Cross & Prusak, 2002; Cummings, 2004; Reagans & McEvily, 2003). Social network analysis (SNA) is becoming an increasingly important paradigm in the organizational literature (Borgatti & Foster, 2003), so much so that some researchers have begun to claim that "the network is the organization" (Contractor, Wasserman, & Faust, 2006). The main difference between SNA and traditional organizational research is that traditional research is concerned with the characteristics of entities, whereas SNA is primarily concerned with the relationships between those entities. SNA examines the structure of relationships in a network, predicting outcomes for the network that

result from the features of that particular structure. Two building blocks comprise social networks - nodes and ties. Nodes are any entity in the network (e.g. people, companies, information systems), whereas ties describe the relationships between those nodes.

Typical examinations of social networks study the structural patterns of the network along a single mode of interactions – individuals interacting with other individuals, workgroups interacting with other workgroups, companies interacting with other companies. Social network analysis, however, is also capable of examining more complex networks comprised of different types of nodes (Wasserman & Faust, 1994). A classic but under-used network conceptualization is known as the two-mode or affiliation network (Borgatti & Everett, 1997; Faust, 1997). A two-mode network is a general network structure comprised of two fundamentally different types of entities that cannot be examined equivalently with one another. An affiliation network refers to a particular use of a two mode network to study how people are related to each other through membership in shared events (Carrington, Scott, & Wasserman, 2005). In other words, affiliation networks are a special case of two-mode network typically comprised of two particular types of nodes – people and events.

The classic example of two-mode affiliation networks in the social network literature here is Davis's study of the social events of Southern women (Wasserman et al., 1994). Davis used a two-mode network to identify how the events created a social structure among the women and how the women created commonalities among the events. Women attend events with a small group of other women, and certain events tended to attract certain types of attendees. Events were shown to be socially similar if they attracted a similar cadre of attendees, and women were shown to be socially related it they attended a large number of events together. Two-mode affiliation networks have also been used to study corporate board interlocks, political

affiliations among legislators, opinions of Supreme court justices, scientific communities, and many other types of clustered social organizations (Carrington et al., 2005).

The Social Structure of Collaboration

Researchers frequently use SNA to study collaborative environments because the relationships that connect people to one another in a group can serve as conduits for knowledge and information flows among that group of people. The structure of those relationships and a particular node's position within that structure, therefore, will have a significant bearing on the overall characteristics of the collaborative environment. The earliest studies of information and knowledge sharing in social networks examined how information in chain letters flowed through a network of people (Milgram, 1967; Travers & Milgram, 1969) and how the types of relationships maintained by job seekers influenced their ability to obtain information to find jobs (Granovetter, 1973). More recent work has shown that the same network structures and knowledge-sharing features identified in real-world interpersonal networks also extend to the virtual world of online collaboration (Ahuja & Carley, 1999; Ahuja, Galletta, & Carley, 2003).

Social network analysis can be used to study the collaborative relationships that occur on Wikipedia. If editors of an article are also actively contributing to other articles, they can transfer information and knowledge from one collaborative environment (i.e. article) to another. The information and knowledge contributed to one article may be transferred to benefit the development of another. Editors may transfer *content knowledge* that is directly relevant for the content development of a related article. For instance, information included in an article on Autism might be relevant for the development of an article on Asperger's Syndrome, a different but related condition. Thus, if editors of the article on Aspberger's syndrome have access to the

information in that article on Autism, these editors can use this content knowledge to improve the former article with information from the latter (and vice-versa). This content knowledge may also involve an awareness of the references that can be sought for relevant information and cited to bolster its credibility.

Editors may also transfer *process knowledge* in which the individual learns how to effectively collaborate with others in a variety of settings and conditions. The greater variety of collaborative instances in which an article's editors have been involved will increase the amount of experience they have to draw upon to guide their own collaborative processes. Collaborators may learn how to effectively resolve conflict on a particular issue, for instance, because they had observed a similar issue arise in the development of another article. Process knowledge might also capture an editor's sense for effective presentation style. Researchers have found that effective packaging of information and knowledge is often associated with perceived quality (Markus, 2001). As editors are involved in a number of different articles, information will be presented in various different ways and different members of the community will have different reactions to the presentation of any given article. The more exposure editors have to different presentation styles; he or she will gain a better sense of effective and ineffective ways to present a variety of types of information. This sense can be used to present information more effectively on the articles in which the editor is directly related.

Editors might also transfer *reputational knowledge* about which other collaborators contribute most to (or detract most from) the effective development of knowledge objects. If one contributor has a reputation for high-quality work on other articles, editors may be more willing to trust their insight on a controversial issue than another editor who is relatively unknown. Conversely, if another editor is known for disruptive behavior or particular types of inappropriate

editorial activity (e.g. pushing a particular fringe point of view), this behavior might be identified early and prevented from having a deleterious effect upon the collaborative environment.

In short, different forms of knowledge gained by working on one article can be applied to the development of other articles in which the editor is involved. Thus, in order to understand article quality, it may not be sufficient to explore the direct collaborative activity related to a particular article without also exploring more generally where an article is situated within the overall collaborative structure of knowledge flows on Wikipedia. This paper uses a two-mode affiliation network to study the social structure that exists as the result of editors collaborating across multiple articles on Wikipedia. This social structure will determine the flow of information and knowledge to a particular article, which will in turn affect article quality. Articles whose editorial community can access more content, process, and reputational knowledge present on Wikipedia are more likely to be higher quality than articles that cannot access this knowledge.

Using Centrality to Measure Collaboration

A measure that is frequently used to understand the position of an individual node within knowledge flows in a collaborative environment is *centrality*. Centrality refers to whether an individual node is situated within the core or the periphery of a social network (Brass, 1995; Scott, 2000). Centrality in collaborative networks is often positively associated with the benefits an individual extracts from a collaborative network. Individuals who are central in a collaboration network are more likely to be involved in innovative activity (Ibarra, 1993), typically enjoy higher satisfaction with the collaborative environment (Baldwin, Bedell, & Johnson, 1997), typically possess knowledge that is desired by others in the network (Perry-

Smith, 2006; Perry-Smith & Shalley, 2003) and often perform better as the result of their ability to access information more quickly and easily as a result of their central position in the collaboration network (Mehra, Dixon, Brass, & Robertson, 2006; Sparrowe, Liden, Wayne, & Kraimer, 2001). Centrality has also been shown to be a particularly critical measure in online social interactions. The centrality of project managers in an open-source developers social network has been shown to be related to the success of the projects in which they are involved (Grewal, Lilien, & Mallapragada, 2006). Centrality in online collaborative environments are associated with an individual's ability to exchange knowledge with others in the collaborative environment (Wasko & Faraj, 2005) and it is positively associated performance on tasks that are associated with the exchange of that information (Ahuja et al., 1999; Ahuja et al., 2003).

Although previous research has typically examined the influence of relational structure on individuals in interpersonal networks, the importance of centrality holds at other levels as well. For instance, workgroups that are more central in a network of other workgroups tend to be able to leverage similar information advantages as has been identified at the individual level (Tsai, 2001). Groups that have ties with other groups are able to use these relationships to transfer knowledge and, in turn, improve the performance of the entire group (Cummings, 2004; Hansen, 1999, 2002). Similarly, firms that are central in a network of strategic alliances are more likely to benefits from those strategic alliances than firms that are peripheral in that network of alliances (Owen-Smith & Powell, 2004; Walker, Kogut, & Shan, 1997).

Network researchers have developed a number of different measures of centrality – eigenvector, degree, betweenness, closeness, to name only a few (Wasserman et al., 1994). Although these measures are somewhat correlated with one another, depending on the size of the network and the measures used, they often capture important differences about the structure of

relationships in a network. The particular measure of centrality must therefore be chosen carefully in order to most accurately represent the type of network dynamic of interest (Borgatti, 2005). Both the local and global structure of the social network can influence how an individual node benefits from the information flows in the network. Local structure reflects the benefits a node receives from the nodes to which it is directly connected, whereas global structure reflects the benefits a node receives both directly and indirectly from all of the nodes in a network. Two measures of centrality are often used in conjunction with one another to capture the features of the local and the global social network in combination with one another – degree centrality and eigenvector centrality (Faust, 1997; Friedkin, 1991). Degree centrality captures the centrality of a node in the local network by calculating the number and strength of the relationships of which that node is involved. Eigenvector centrality captures the centrality of a node in the global network by weighting the value of a node's direct ties as a function of all of the relationships those nodes are involved.

Degree Centrality as a Measure of the Local Network: The local network structure suggests that the types of relationships directly maintained with other nodes in the network influences the type of information an individual can obtain from its social network. Relationships in local networks are often evaluated in terms of number (the total number of direct relationships maintained) and strength (frequency and depth) of interaction. The number of ties maintained by a node is often positively related to information flows. For instance, the greater number of ties maintained by a node is more likely to result in the successful response to a query (Constant, Sproull, & Kiesler, 1996), and more ties are positively related to upward mobility in the workplace (Podolny & Baron, 1997).

The strength of ties in a node's local network has also been shown to be related to the flow of information and knowledge, but in somewhat paradoxical ways. Weak ties are better for searching for information in a network, but strong ties are better for transferring information (Hansen, 1999). Further research demonstrated that the relationships of tie strength to knowledge transfer are also dependent on the type of information transferred. Weak ties have been shown to be better for transferring explicit knowledge, but strong ties are better able for transferring tacit knowledge (Hansen, 2002). Most of the paradoxical benefits related to tie strength, however, are related to the inherent limitations of people to maintain a limited number of ties (Burt, 1997; Hansen, 1999). A two-mode incidence network, like the one used here, is not subject to these same limitations. Previous research has found that large numbers of editors do not appreciably interfere with the ability of a group to collaborate with one another to develop Wikipedia articles (Kane and Fichman forthcoming). Therefore, degree centrality that measures the number and strength of a node's local network is expected to have a uniformly positive effect on article quality.

Degree centrality captures how well an article is directly related to other articles on Wikipedia. It captures the capacity with which information and knowledge can be transferred from other collaborative environments to a certain article through editors who collaborate on multiple articles. High degree centrality suggests that an article's connection with other articles is both greater in number and stronger. An article with high degree centrality is more likely to have editors that are both able to identify needed knowledge within other articles and transfer that information to improve the article. Low degree centrality suggests that an article's connection with other articles is both fewer in number and weaker. This network structure makes it less likely that an article will benefit from information and knowledge contained in

other articles in the network. The more information and knowledge that is available to be transferred to an article; the higher quality that article is likely to be.

H1: The degree centrality of a given article in the two-mode incidence matrix of articles and editors will be positively related to article quality.

Eigenvector Centrality as a Measure of the Global Network: Eigenvector centrality captures the article's position in the global network structure – all of the nodes and ties that comprise a network. The global network structure suggests that the network relationships beyond the direct relationships can also be associated with the information and knowledge available to a node. The argument is simply an extension of the argument related to the effect of the local network to the entire network. If the amount of information and knowledge available to a node is related to its local network structure, then the local network structures of all the nodes in one's own local network will also be related to the amount of information and knowledge that can be transferred from that node. Eigenvector centrality has often been called the "total effects" centrality, because it captures the centrality of a given node by assessing the centralities of the nodes with which it is connected (Faust, 1997; Friedkin, 1991). Eigenvector centrality weights the value of the nodes in one's local network as a function of how well they are connected to other articles and, iteratively, how well those third-degree articles are connected to other nodes, until the centrality of every node in the network is a function of the centrality of all other nodes in the network (Bonacich, 1972).

Eigenvector centrality may be illustrated through simple examples. An individual who only has three friends may not seem particularly important, until one learns that these three friends are Warren Buffett, The Pope, and the President of the United States. Relationships with these people would make one more important by extension. Although the number of direct

connections may be relatively limited, the network that one can access through those three friends is quite extensive and influential. Another example of eigenvector centrality is Google's Page Rank Algorithm (Brin & Page, 1998). Google uses a form of eigenvector centrality as a core component of its page rank algorithm – web pages are more important if they are linked to by important websites (e.g. CNN, ESPN, Microsoft) than if they are linked to by less important websites (e.g. a college professor's personal webpage). Google is now adopting similar approaches to determine the rank of people in online social networks, such as Facebook, MySpace, and LinkedIn.

Eigenvector centrality evaluates the likely impact of knowledge transfer from the nodes in an article's local network as a function of the information and knowledge available to those nodes through their local network. Collaborative environments that are well connected to other active collaborative environments are better able to access process, content, and reputational knowledge that can be of use to article development for the other articles in their network than is an article that has a relatively small and isolated collaborative environment. Articles that are heavily edited and well connected to other well-connected articles are more likely to contain more information and knowledge from which others can benefit. Degree centrality would treat these articles as equivalent sources of knowledge, but eigenvector centrality would capture the fact that that certain article are likely to be a better source of information and knowledge than other articles as a result of its position in the network.

H2: The eigenvector centrality of a given article in the two-mode incidence matrix of articles and editors will be positively related to quality.

Research Setting and Method

Drawn from the Hawaiian word meaning "quick," a wiki is simply a website that anyone can edit. Wikipedia uses a wiki platform to host an open-source encyclopedia. Established in 2001, the English version of Wikipedia has, as of this writing, developed over 2.5M separate articles. The 40 other languages in which Wikipedia is published also hosts an additional eight million articles. Anyone can make an edit to any article on Wikipedia. When the user does so, the editor's identity, the changes he or she made to the article, a description of the change, and the time of the change are all recorded by the system. Other users can be automatically be notified of any changes made to a particular page, and they can also undo any set of edits to those of a previous version.

Wikipedia runs on an open source Wiki platform known as MediaWiki that consists of a number of features. First, this article contains the text of the article, multimedia images of the subject matter at hand, links to other relevant information within Wikipedia, and a reference list to support the facts contained in the article. Second, two tabs on this article page give access to the community interaction which yielded the article at hand. The "history" tab displays data on every edit that has been made to the article, including a summary of the edit, whether it is a minor edit, the time of the edit, and the user who made the edit. The "discussion" tab displays a forum provided on which editors can consider the merits of possible additions, deletions, or edits, can resolve differences, or jointly decide on other issues related to the development of the article. In addition to the article, history, and discussion tabs for each article, a similar information structure exists for every Wikipedia user. The user page displays information about that user, the discussion pages host interactions between that user and others on Wikipedia, and

the history page reveals all edits to each of the pages. Although the user page is dedicated to a particular user, any member of Wikipedia can still edit that page.

This paper focuses on a relatively small subset of articles on Wikipedia to enable the collection of rich and social network data. A random sample of articles, however, would not likely yield the social network features in which this paper is interested, so this paper focuses on a small subset of articles contained in a single Wikiproject. A Wikiproject is a group of editors dedicated to develop, maintain, and organize articles related to a particular topic. There are literally hundreds of Wikiprojects on Wikipedia, dedicated to a wide range of topics, ranging from the mainstream to the obscure. Because of their common subject matter, Wikiprojects are more likely to share editors with one another, resulting in a smaller and more clearly defined sub-network of articles and editors than would be reflected in a wider sample of Wikipedia articles. The smaller sub-network also allows the researcher to conduct qualitative analysis on the top editors and their editorial activity to understand the nature of the interactions and the collaborative environments better. Further, the Wikiprojects usually rate the quality of each article under its purview on a 5-point quality scale, providing the researcher with the ability to separate full articles from various types of proto-articles such as article stubs (articles in early stages of formation) or lists. Wikiprojects also provide a certain experimental control for the subject of the articles. Articles related to medicine may be evaluated against entirely different standards for quality than structure than articles on pop-culture, and studying the articles dedicated to a particular Wikiproject limits the impact of potentially confounding factors.

Medicine Wikiproject: This project studies a selection of articles from the Medicine Wikiproject, which was chosen for a number of reasons. First, health information has always been one of the key uses of online information (Ferguson & Frydman, 2004). As such, the

information contained in these articles will be highly relevant to the general population, attracting a wide audience of interested readers. To test this assessment, a Google search for 300 health terms was conducted, finding that the Wikipedia article was listed as the first result 60% of the time and in the top three results 90% of the time. Thus, medical information on Wikipedia is likely to be used heavily by the general public, ensuring that these articles will likely receive widespread editorial and popular attention.

Second, the Medicine Wikiproject has developed a set of guidelines for developing medical articles on Wikipedia that are stricter than a typical Wikipedia article. For instance, only peer-review academic articles are considered authoritative references (in contrast to high-quality articles in the popular press that are acceptable elsewhere on Wikipedia), and official scientific terms must be used instead of colloquial expressions (e.g. Myocardial Infarction instead of heart attack). Given the additional criteria, the somewhat restricted access of most peer-reviewed sources to the general public, and the required knowledge of accepted medical terminology, the Medicine Wikiproject represents a distinct, insular sub-community on Wikipedia. This feature of the Medicine Wikiproject will permit better isolation of the social network features, allowing this paper to collect data that is both contextually rich but computationally adequate for the research questions in which this paper is interested.

Third, the Wikiproject also assigns an "importance" rating for each article, depending on the relevance of the medical topic to the general public. Cancer, for instance, is a medical topic that is relevant to a wide population of patients, caregivers, and medical professionals. It is likely to involve substantively different collaborative dynamics than a condition such as Leopard Syndrome, a relatively rare genetic disorder. This enables the researcher to control for the subject matter of a given article. More important topics are likely to be the subject of

qualitatively different types of editorial activity than less important topics. Limiting this study to articles only in the Medicine Wikiproject enables us to control for these collaborative differences that may exist because of differences in topic importance that would not be possible in a more general study of Wikipedia articles.

Data Collection:

For the purposes of this study, the researcher selected a sample of articles that would create a relatively balanced distribution of articles in terms of quality and importance, while also gathering enough data for adequate statistical power. The Medicine Wikiproject rates the quality of articles on a 5 point scale: Featured Article (best), A-quality, Good Article, B-quality, Start-quality (worst). The researcher first selected all of the articles in the Medicine Wikiproject that has been granted Featured Article, A-quality, or Good Article Status by the broader editors of Wikipedia. This resulted in a sample of approximately 100 articles. Then, a separate set of 200 articles was selected at random from the approximately 7000 articles that had been labeled B-quality and Start quality by the Wikiproject. The result was an overall sample of 300 articles for analysis. Such traditional sampling methods often introduce the capacity for severe bias in network research, because dyadic measures are not independent of one another (Wasserman et al., 1994). Nevertheless, researchers have noted that two-mode networks are not dyadic and, thus, are not subject to many of the assumptions that govern the analysis of one-mode networks (Faust, 1997). As a robustness check, the researcher divided the data into two random, equally-sized subsets of the data, deriving each of the network statistics used and conducting all analyses described in this paper on these subsets of network data. In all cases, results are qualitatively the same as those presented throughout the paper, providing strong evidence that the sampling methods adopted by this paper did not bias the results presented here.

To minimize the impact of article variations over time, data was collected over two weeks during the 1Q of 2008. A team of research assistants collected the editorial history of each of the 300 articles selected. This was accomplished by examining a page history statistics using a tool provided by Wikipedia (Kane and Fichman 2008). The reliability of this tool was established by comparing the results of the tool to the raw data collected and analyzed directly from the article history, and it was found to accurately reflect the editorial history preserved by Wikipedia. The researcher also examined the user pages of the top 10% of editors to learn about who was editing the medical articles. Approximately 1/3 of the editors claimed some sort of professional medical experience, another 1/3 were patients with particular conditions, with the remaining 1/3 having no explicitly stated connection to medicine or medical articles. Further, the researcher submitted a survey to the most prolific editors on the Wikiproject in an attempt to understand this particular collaborative environment better. Many of the details found on the editor pages were corroborated using various details provided by these editors. Some of the editors were medical professionals (often European), were patients with the conditions they edited on, or were simply general Wikipedia enthusiasts.

Respondents indicated that the Medicine Wikiproject was, in fact, a largely independent sub-community on Wikipedia. An examination of the editorial history of top editors demonstrated that approximately 75% of the most active editors of medical articles contributed virtually exclusively to medical articles, electing not to edit other types of articles on Wikipedia. Editors also exhibited considerably different patterns of editing. Some editors focused very heavily on a few related articles, whereas other editors contributed broadly across the entire Medicine Wikiproject. The top two editors in terms of total number of edits represented a good contrast of these editorial styles. Although each had made roughly the same number of edits, one

had actively contributed to over 150 articles whereas the other had focused almost his entire editorial activity on a single article. These different editorial patterns will be captured by the social network measures used here but would have been overlooked in previous studies that examined articles independently. Some editors will be more effective conduits of information between articles than others.

Dependent Variables:

The Medicine Wikiproject evaluates the articles on a 5-point scale (Featured, A, Good, B, and Start). Two of these categories, Featured and Good articles, are assigned by Wikipedia in general. Although any member of Wikipedia may nominate an article for consideration as a Featured or Good Article, the designation of Featured Articles and Good Articles are awarded only after the article is subject to a rigorous peer review process. See Appendix A for Featured Article Criteria. Once nominated, Wikipedians are invited to consider the merits of the article and vote on whether the article should be granted Featured/Good article status. Wikipedians may vote to either support the nomination or oppose the nomination, citing the reason for their vote. Approximately 0.01 % of articles on Wikipedia have achieved Featured Article Status (2200 of 2.5M, as of this writing), and approximately 0.05% of articles have achieved Good Article Status (5,000 of 2.5M). The other three categories (A, B, and Start) are assigned by a consensus of members of the Wikiproject. A-quality articles were often also regarded as good articles by Wikipedia in general, and the Wikiproject decided to provide them extra recognition. Neither B-quality nor Start-quality articles had received any formal recognition for quality by Wikipedia generally. Several had been nominated, but the peer review process had deemed them not high quality enough to merit recognition as either a Featured or Good Article.

The validity and reliability of this rating system as a proxy for article quality has been established by and used in previous research (Kittur et al., 2008). We also independently employed a team of graduate nursing students to evaluate the content of a small subset of 20 articles in our sample, who confirmed the reliability of this rating system. It should be noted that the dependent variable used here does not determine the absolute quality of an article, but evaluates its relative quality to other articles in the sample. Further, a battery of robustness checks was conducted on these dependent variables to ensure that results were not an artifact of the dependent variable construction. Various combinations of category groupings were tested (e.g. grouping Featured and Good Articles and/or Start- and B-quality articles into a single category), and the data associated with particular categories were also dropped (e.g. A-quality and/or Start-quality). Appropriate regression models were used to examine the different ways the dependent variable was operationalized (e.g. dichotomous or ordinal). None of these robustness checks affected the overall results of the model presented here, so these findings are robust with respect to the particular form of the dependent variable chosen. Since the form of dependent variable did not affect model results nor did the particular sample chosen, this paper uses the full 5-point rating in order use all of the data available for analysis.

Independent Variables:

The identity and volume of editorial activity of all editors who had contributed at least 5 edits to each of the articles was obtained, yielding a total population of 1800 editors and the total number of edits editor made to each article in the sample. This data was used to construct an edgelist, which was then imported into UCINet 6.198 to construct the 300x1800 2-mode matrix of articles and editors. UCINet was used to multiply the 2-mode matrix by its transpose, creating the 300x300 incidence matrix of articles and the shared editorial activity between them. A visual

representation of the resulting 2-mode network constructed using NetDraw 2.081 can be found in Figure 1. This visualization was constructed using an iterative process that positions each node in the network based on their relationships to all other nodes using a multidimensional scaling-like algorithm (Huisman & Duijn, 2005).

One of the most common approaches to analyzing two-mode networks is to construct an incidence matrix that is constructed by multiplying the two-mode network by its transpose (Faust, 1997). This method creates a new social structure by examining one mode of the network from the perspective of the other, treating one mode a node in the incidence matrix and the other mode as the ties which join them together. In other words, the incidence matrix can examine how events are related to each other as a result of its shared members or how members are related to each other by their participation in common events (depending on the order of the matrix multiplication). For instance, Davis showed that certain types of events were similar to one another because they shared a similar list of attendees (Wasserman et al., 1994). Since the purpose of this paper is to examine how editors may transfer information and knowledge from one article to another by their participation in both collaborative environments, the incidence matrix is constructed in such a way that the articles are treated as nodes and the editors are the ties which join them together. This approach conceptualizes editors as conduits (i.e. ties) for transferring knowledge and information from one article (i.e. node) to another.

The independent variables of interest were then derived from the incidence matrix constructed using this data. First, independent variables were created from the two-mode and incidence matrix created in UCINet 6.198. There are two available methods for constructing incidence matrices, one that involves computing the cross products of the two matrices and the other more conservative measure that involves taking the minimums of the two matrices. Both

methods were employed and the network variables were constructed on each. Model results were qualitatively the same for both sets of network variables. The cross-product method was chosen for final analysis, because it minimized the correlation between the network variables. *Degree centrality* was calculated in UCINet 6.198, and calculated the total number and strength of ties maintained by a given node. *Eigenvector centrality* was also calculated in UCINet 6.198 on the incidence matrix. Bonacich and Lloyd (2004) define eigenvector centrality as follows:

> Let $\boldsymbol{A}$ be a symmetric adjacency matrix, where $aij = aji = 1$ if i and j are connected in a network and $aij = aji = 0$ otherwise. The eigenvector measure of centrality x is the solution to the following matrix equation: $\boldsymbol{A}x = \lambda x$ *(p. 331-332).*

Eigenvector centrality calculates a node's centrality in the network as a function of which other nodes in the network is connected to and how central those nodes are in the network. It captures the overall prominence of the articles with which the focal article shares editors and editorial activity.

Control Variables: Since factors other than the social network may affect an article's rating, several control variables were also used. First, information was collected regarding the direct collaborative activity related to each article. It may be that the time, energy, and effort dedicated to the development of each article may affect quality without the need for editors to transfer information from other articles. Previous research has found that the number of edits and editors are positively related to article quality (Wilkinson et al., 2007). Unfortunately, these two measures are highly correlated with one another, introducing severe multicollinearity into the analytical models. To account for these characteristics of the data, the models use *total number of editors* and then *average number of edits per editor*. Measures of the concentration of collaborative activity have been found to be positively related to article quality (Kittur et al., 2008), so this revised measure seems consistent with previous research. These measures were

taken separately for editorial activity on the article and for editorial activity on the discussion page dedicate to the article. Both types of collaboration may be independently important for understanding article quality, so these aspects of collaborative activity are collected independently from one another. Thus, using four measures are used to capture direct collaborative activity on an article – *number of unique editors (article), number of editors (discussion), average edits/editor (article), average edits/editor (discussion).*

Second, the models control for the relative popularity of particular articles. An article may be higher quality if it attracts more attention. Studies have shown that editors may also be more likely to contribute to a high-traffic or high-profile article (Zhang et al., 2007). *Article Traffic* is used to measure how often an article is viewed in terms of number of page hits. Since "anyone can edit Wikipedia," a larger number of people who read a given page are more likely to notice and correct any factual errors or contribute missing knowledge. Traffic statistics were gathered from Wikipedia, using data from Alexa.com. Traffic is operationalized by the total number of hits to a page during the first quarter of 2008. Studies have also shown that patients most frequently click on only the top search result when searching for health information, spending little time exploring lower search results (Eysenbach & Köhler, 2002). Articles that have the top Google search rating, therefore, may be of higher quality than articles ranked lower as a result. *Google Rank* is a binary variable that captures whether the article is the top Google search result for a search for that medical term.

Third, the models control for the topic of the article. Articles on more important topics are also likely to receive more attention than other articles. For instance, it is more likely that people will search for information on a relatively common or a more severe illness or condition (e.g. tuberculosis or cancer) than they will on less common or severe conditions (e.g. Leopard

syndrome or Flynn effect). It is also likely that people are more likely to have knowledge about common conditions than about relatively rare conditions. Thus, the base of knowledge is likely to be greater about topics that are of a higher importance. Articles with highly important topics might also attract a greater number of divergent opinions about what information should be contained in the article. Extended controversy and conflict can detract from the overall quality rating of the article (Jehn, 1997). *Topic importance* captures the significance of an article's subject matter to the general population. The Medicine Wikiproject rates the importance of a given topic on a 4-point scale (4 being top importance, and 1 being low importance).

Finally, editorial anonymity may affect the nature of the social relationships, so we control for the influence of anonymity on the collaborative environment. When an editor makes an edit to the discussion or article page, his or her identity is automatically recorded in the system. If the editor is not logged into the system, the editor's identity is recorded as an anonymous IP address. All edits made during that session are recorded to the same IP address. The raw number of anonymous editors was too highly correlated with total number of editors to include in the same model, so the *percentage of anonymous editors* is used. This variable was calculated as the total number of anonymous editors divided by the total number of editors on an article or discussion page. Because previous research has hypothesized different effects for anonymity depending on whether it is related to contributions or group process, anonymity on the talk pages and the article pages are controlled for separately (Connolly, Jessup, & Valacich, 1990; Jessup, Connolly, & Galegher, 1990; Sia, Tan, & Wei, 2002).

Data Analysis:

First, an examination of the correlation matrix (Table 1) demonstrated some relatively high correlations between certain variables that were to be expected (e.g. between centrality measures). The model was analyzed, therefore, to ensure multicollinearity would not pose a problem for the data analysis. The final model does actually show a moderate degree of multicollinearity, two of the variance inflation factors (VIF) in the final model being slightly elevated (4.5 and 5.5). The variables causing the high multicollinearity, however, were the total number of editors for the article and discussion pages. Although some statisticians suggest that VIFs below 10 are acceptable (Neter, Wasserman, & Kutner, 1990), others adopt a more conservative number of around 5. The model was tested to ensure that the multicollinearity did not affect final results. The variables associated with the high VIFs were dropped from the model, both individually and in combination. When these variables are removed from the model, the VIFs for all models drop to below 3 without any appreciable effect on the variables of interest and all model results were the same with either or without both of these variables. Since the mulicollinearity was only moderately high and it did not affect the model results, the full model was retained.

Data was then analyzed using the ordinal regression function in SPSS 16. Ordinal regression is used when there is a progressive relationship within a categorical dependent variable, but it is unclear the magnitude of different between the categories. For instance, the observer may know in the Olympics which athletes have won the gold, silver, and bronze medals without knowing the final scores of any of the athletes. Ordinal regression is a generalized linear model of the following form:

$$\text{link}(\gamma_{ij})=\theta_j-[\beta_1 x_{i1}+\beta_2 x_{i2}+\ldots+\beta_p x_{ij}]$$

link( ) is the link function
γ_{ij} is the cumulative probability of the jth category for the ith case
θ_j is the threshold for the jth category
p is the number of regression coefficients
$x_{i1}\ldots x_{ip}$ are the values of the predictors for the ith case
$b_1\ldots b_p$ are regression coefficients

Various link functions can be chosen based on the distribution of the underlying variables. A negative log-log function was chosen because there was a higher density of Start- and B-quality articles in the sample than Featured or Good Articles. Nevertheless other link functions were chosen used as a robustness check, and results are robust with regard to the choice of link function. The model fitting statistic was highly significant and the Pearson lack-of-fit statistic was insignificant, suggesting that the resulting model fit our data fairly well. Further, both the Nagelkerke and Cox-and-Snell R^2 statistics of the final model were around .50, suggesting that our model explained a considerable amount of variance in article quality.

Results

The full results can be found in Table 2 at the end of the paper. Hypothesis 1 stated that an article's degree centrality, representative of how closely a particular article was related to other articles as a result of their structure of shared editors, was highly significant in the expected direction (Wald = 7.29, $p < .01$). This result suggests that the degree to which articles can benefit from the knowledge contained in the collaborative processes that occur in relation to other articles on Wikipedia will be positively related to article quality. Editors do appear to transfer information and knowledge from one collaborative effort to others in which they are involved.

Hypothesis 2 stated that eigenvector centrality, capturing how well connected an article's alters are connected to other articles as a result of their structure of shared editors, would be positively related to article quality. Hypothesis 2 was highly significant and in the expected direction (Wald = 12.60, $p < .001$). Articles that are more central in the overall collaborative structure as a result of being connected to other well-connected articles are more likely to be of higher quality. Connections with articles that are more likely to be a valuable source of information and knowledge are more valuable for article quality than connections with articles that are relatively obscure and unlikely to be a source of information and knowledge.

Although not formally stated as hypotheses, it is also worth noting several results related to the control variables. First, the topic importance variable is the most highly significant variable in the model (Wald = 53.72, $p = .001$). Surprisingly, however, it shows a strong *negative* relationship between article quality and topic importance. At one level, this result seems counter-intuitive. If Wikipedia is an environment motivated by altruism and pro-social behavior, as has been recognized in previous studies of online collaboration (Constant, Kiesler, & Sproull, 1994; Constant et al., 1996; Wasko & Faraj, 2000), it would seem that more important topics would attract a higher number of well-meaning editors that would improve the overall quality of collaboration. The data demonstrates that this is clearly not the case. An alternative explanation is that highly important topics receive a greater deal of attention and scrutiny, bringing a wider variety of individuals with diverse knowledge. These differences may be related to greater controversies and conflict in article development, leading to sub-optimal collaboration. This interpretation is partially supported by the correlation matrix, which shows that although there is no significant correlation between total edits on an article and the importance of the article, there is a positive correlation between the number of editors and

importance. It is further interesting to note that topic importance is not significantly correlated to either of the measures of centrality on the incidence matrix, suggesting that the difference does not result from the position of important articles within the overall collaborative structure. Regardless, it is worthwhile to note the finding here that more important articles on Wikipedia tend to be of lower quality and the higher quality articles tend to be dedicated to less important topics.

Second, consistent with previous literature (Connolly et al., 1990; Jessup et al., 1990; Sia et al., 2002), anonymity had differential effects, depending on whether it occurred on the discussion pages or the article. The percentage of anonymous editors is significant and positively related to article quality when it occurs on the article pages (Wald = 10.72, $p<.001$), and significant and negatively correlated to article quality when it occurs on the discussion pages (Wald = 3.71, $p<.05$). These findings provide further evidence that anonymity has a somewhat paradoxical and differential effect on article quality – positively related in certain instances but negatively related in others.

Third, very few of the variables related to direct collaboration were positively related to quality. Only the average number of edits per editor on the article pages was found to be positively related to article quality in the final model. This suggests that the amount of direct collaborative activity could actually have little influence on article quality. Further, it is interesting to note that total number of editors for the article is significant in early models, but it becomes insignificant when social network measures are added. These results underscore the importance of considering the connections between articles when examining the collaborative structure of Wikipedia. Failing to do so may not only overlook a key variable associated with

article quality, since the models where the social network measures are included explain considerably more variance, but it also may provide biased and/or erroneous results.

Discussion

Conducting a study of the collaborative processes directly related to and occurring between 300 articles about health and medicine-related topics on Wikipedia, this study sought to explore whether certain collaborative features are related to article quality. In general, this study found generally good support for its hypotheses. By and large, the social network measures of collaboration between articles offered considerable insight into article quality and more so than did the control measures of direct collaboration.

Theoretical Implications: This study makes a number of important contributions to theory. First, it establishes that certain collaborative features are related to overall article quality on Wikipedia. Although Wikipedia editors are fond of noting that their collaborative efforts "work in practice, not in theory," this study suggests that there are certain features of Wikipedia's distinctive brand of online collaboration that can be explained by theory. This study should only be regarded as a first step in identifying these effective practices. Future researchers can identify and establish other features that may be related to article quality.

Second, this study establishes the importance of understanding the collaborative structure between articles to gain the most complete picture of collaboration on Wikipedia. Most previous research has viewed Wikipedia as a collection of independent articles, whereas this study has examined the collaborative structure that results from multiple editors collaborating on multiple articles on a single platform. It finds that this collaborative structure between articles is critically important for understanding article quality. Editors serve as conduits for transferring

information and knowledge from the collaborative activity of one article to other articles under development. How editors work on multiple articles simultaneously may be a critically important feature of collaboration using Internet-based collaborative tools and should not be overlooked. Not only the direct connections matter, as measured by degree centrality, but so does an article's position within the wider network of articles to which it is not directly connected.

Managerial Implications: This study also has a number of implications for management practice. Many companies are seeking to use wikis and Internet-based collaborative tools in their companies to support collaboration within, across, or even outside organizational boundaries (Howe, 2008; Li et al., 2008; Tapscott et al., 2006). Employees can use these tools to be more productive and companies often create customer communities that can collaborate with one another to generate important insight for the company and each other (Wagner et al., 2006). Although Wikipedia represents a unique collaborative environment, making it important to carefully qualify any generalization of the results presented here, companies can learn important insights from this study to guide their own collaborative efforts.

This article identifies which collaborative features are most likely to be associated with collaborative success. Namely, editors being active across multiple collaborative environments and judicious use of anonymity in these collaborative environments are all positively associated with collaborative success. How might these findings inform collaboration in other environments? First, the structure of collaborative activities can be shaped by the generative rules under which the collaboration operates (Kogut, 2000), particularly the incentive structure in the case of intra-organizational collaboration (Ba, Stallaert, & Whinston, 2001). These results suggest that managers would want to cultivate a balanced collaborative environment.

Individuals should become invested in the collaborative environments in which they are engaged resulting in a high concentration of editorial activity (Kittur and Kraut 2008), but individuals should also be encouraged to be involved in multiple collaborative environments to enable them to transfer information and knowledge between collaborative environments. Neither an overly narrow effort that focused considerable attention on a single article nor and overly broad effort that dabbled in many articles would appear as effective as this balanced approach.

Limitations and Future Research: This study has a number of limitations that create possibilities for further research. First, choosing to study a particular subset of Wikipedia helped this paper isolate the phenomenon of interest (e.g. the social network structure of editors) and control for a number of other article characteristics (e.g. article importance) that would not have been possible otherwise. This approach, however, does limit the degree to which our findings can be generalized to other settings either on Wikipedia or in collaborative environments using other Internet-based collaborative tools. Although there is no pressing reason to think that the findings presented here are necessarily dependent on the research setting, the findings presented here should be generalized carefully. Future research should seek to replicate these findings in other settings to confirm the importance of the network structure between articles.

Second, this study focuses on a snapshot of the networks and article quality at a single moment in time. Although a common limitation of social network studies, these collaborative networks do change and evolve over time. This study offers little to explain how particular network structures develop over time, how the dynamic nature of networks and article quality are related to one another, and how these influences can be affected. Network dynamics is an important frontier for social network research and the collaborative structures that occur between editors and articles would be a rich setting for future research to examine these articles.

Nevertheless, some of the relationships not expressly hypothesized yielded interesting results that merit further inquiry. For instance, the strong negative relationship between topic importance and article quality suggests that articles on more important topics are likely to be of lower quality. This relationship and a further understanding this relationship has significant implications for the value of Wikipedia and other collaborative as an information resource. It would be interesting to explore further what collaborative dynamics yield such a powerful and counter-intuitive result. Further, the divergent effects of anonymity on article quality provide further confirmation that the effects of anonymity in online environments are somewhat task-dependent. Both of these serendipitous findings suggest possible avenues for future research.

Conclusion: Adopting a social network perspective, this paper demonstrates that the collaboration that occurs between collaborative environments may be equally if not more important than the collaboration that occurs within that environment. Implications are that researchers should seriously consider the social structure that occurs across all elements of a collaborative platform, rather than simply those directly related to the particular outcome of interest. Implications for managers point to the need to manage a collaborative platform as a single integrated environment, rather than as a portfolio of independent collaborative efforts.

References:

Ahuja, M. K. & Carley, K. M. 1999. Network structure in virtual organizations. Organization Science, 10(6): 741-757.

Ahuja, M. K., Galletta, D. F., & Carley, K. M. 2003. Individual centrality and performance in virtual R&D groups: An empirical study. Management Science, 49(1): 21-38.

Ba, S. L., Stallaert, J., & Whinston, A. B. 2001. Research commentary: Introducing a third dimension in information systems design the case for incentive alignment. Information Systems Research, 12(3): 225-239.

Baldwin, T. T., Bedell, M. D., & Johnson, J. L. 1997. The social fabric of a team-based MBA Program: Network effects on student satisfaction and performance. Academy of Management Journal, 40(6): 1369-1397.

Benkler, Y. 2006. The wealth of networks : how social production transforms markets and freedom. New Haven [Conn.]: Yale University Press.

Bonacich, P. 1972. Factoring and weighting approaches to status scores and clique identification. Journal of Mathematical Sociology, 2: 113-120.

Bonacich, P. & Lloyd, P. 2004. Calculating Status with Negative Relations. Social Networks, 26: 331-338.

Borgatti, S. & Cross, R. 2003. A relational view of information seeking and learning in social networks. Management Science, 49(4): 432-445.

Borgatti, S. & Foster, P. C. 2003. The Network Paradigm in Organizational Research: A Review and Typology. Journal of Management, 29(6): 991-1013.

Borgatti, S. P. & Everett, M. G. 1997. Network analysis of 2-mode data. Social Networks, 19(3): 243-269.

Borgatti, S. P. 2005. Centrality and Network Flow. Social Networks, 27: 55-71.

Brass, D. J. 1995. A Social Network Perspective on Human Resources Management. Research in Personnel and Human Resources Management, 13: 39-79.

Brin, S. & Page, L. 1998. The Anatomy of a Large-Scale Hypertextual Web Search Engine. Computer Networks and ISDN Systems, 30(1-7): 107-117.

Burt, R. S. 1997. The contingent value of social capital. Administrative Science Quarterly, 42(2): 339-365.

Carrington, P. J., Scott, J., & Wasserman, S. 2005. Models and methods in social network analysis. Cambridge ; New York: Cambridge University Press.

Connolly, T., Jessup, M. L., & Valacich, J. S. 1990. Effects of Anonymity and Evaluative Tone on Idea Generation in Computer Mediated Groups. Management Science, 36(6): 689-703.

Constant, D., Kiesler, S., & Sproull, L. 1994. Whats Mine Is Ours, or Is It - a Study of Attitudes About Information Sharing. Information Systems Research, 5(4): 400-421.

Constant, D., Sproull, L., & Kiesler, S. 1996. The kindness of strangers: The usefulness of electronic weak ties for technical advice. Organization Science, 7(2): 119-135.

Contractor, N. S., Wasserman, S., & Faust, K. 2006. Testing multitheoretical, multilevel hypotheses about organizational networks: An analytic framework and empirical example. Academy of Management Review, 31(3): 681-703.

Cross, R. & Prusak, L. 2002. The people who make organizations go-or stop. Harvard Business Review, 80(6): 104-+.

Cummings, J. N. 2004. Work groups, structural diversity, and knowledge sharing in a global organization. Management Science, 50(3): 352-364.

Devgan, L., Powe, N., Blakey, B., & Makary, M. 2007. Wiki-Surgery? Internal validity of Wikipedia as a medical and surgical reference. Journal of the American College of Surgeons, 205(3, Supplement 1): S76-S77.

Eysenbach, G. & Köhler, C. 2002. How do consumers search for and appraise health information on the world wide web? Qualitative study using focus groups, usability tests, and in-depth interviews. British Medical Journal, 324: 573-577.

Faust, K. 1997. Centrality in affiliation networks. Social Networks, 19(2): 157-191.

Ferguson, T. & Frydman, G. 2004. The first generation of e-patients. British Medical Journal, 328: 1148-1149.

Friedkin, N. 1991. Theoretical foundations for centrality measures. American Journal of Sociology, 6: 1478 - 1504.

Giles, J. 2005. Internet Encyclopedias Go Head-to-Head. Nature, 438: 900-901.

Granovetter, M. 1973. The Strength of Weak Ties. American Journal of Sociology, 78(6): 1360-1380.

Grewal, R., Lilien, G. L., & Mallapragada, G. 2006. Location, location, location: How network embeddedness affects project success in open source systems. Management Science, 52(7): 1043-1056.

Hansen, M. T. 1999. The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits. Administrative Science Quarterly, 44(1): 82-111.

Hansen, M. T. 2002. Knowledge networks: Explaining effective knowledge sharing in multiunit companies. Organization Science, 13(3): 232-248.

Howe, J. 2008. Crowdsourcing : why the power of the crowd is driving the future of business (1st ed.). New York: Crown Business.

Huisman, M. & Duijn, M. A. J. v. 2005. Software for Social Network Analysis. In P. J. Carrington & J. Scott & S. Wasserman (Eds.), Models and Methods in Social Network Analysis. New York Cambridge University Press.

Ibarra, H. 1993. Network Centrality, Power, and Innovation Involvement - Determinants of Technical and Administrative Roles. Academy of Management Journal, 36(3): 471-501.

Jehn, K. A. 1997. Qualitative analysis of conflict types and dimensions in organizational groups. Administrative Science Quarterly, 42(3): 530-557.

Jessup, L. M., Connolly, T., & Galegher, J. 1990. The Effects of Anonymity on Gdss Group-Process with an Idea- Generating Task. MIS Quarterly, 14(3): 313-321.

Kane, G. C. & Fichman, R. G. forthcoming. The Shoemakers Children: Using Wikis for IS Research, Teaching, and Publication. Mis Quarterly.

Kittur, A. & Kraut, R. E. 2008. Harnessing the wisdom of crowds in Wikipedia: Quality through coordination. Paper presented at the CSCW'08: Proceedings of the ACM conference on computer-supported cooperative work., New York.

Kogut, B. 2000. The network as knowledge: Generative rules and the emergence of structure. Strategic Management Journal, 21(3): 405-425.

Lee, G. K. & Cole, R. E. 2003. From a firm-based to a community-based model of knowledge creation: The case of the Linux kernel development. Organization Science, 14(6): 633-649.

Li, C. & Bernoff, J. 2008. Groundswell : winning in a world transformed by social technologies. Boston, Mass.: Harvard Business Press.

Markus, M. L. 2001. Toward a theory of knowledge reuse: Types of knowledge reuse situations and factors in reuse success. Journal of Management Information Systems, 18(1): 57-93.

Mehra, A., Dixon, A. L., Brass, D. J., & Robertson, B. 2006. The social network ties of group leaders: Implications for group performance and leader reputation. Organization Science, 17(1): 64-79.

Milgram, S. 1967. The Small World Problem. Psychology Today, 1(61-67).

Neter, J., Wasserman, W., & Kutner, M. H. 1990. Applied linear statistical models : regression, analysis of variance, and experimental designs (3rd ed.). Homewood, IL: Irwin.

Owen-Smith, J. & Powell, W. W. 2004. Knowledge networks as channels and conduits: The effects of spillovers in the Boston biotechnology community. Organization Science, 15(1): 5-21.

Perry-Smith, J. E. & Shalley, C. E. 2003. The social side of creativity: A static and dynamic social network perspective. Academy of Management Review, 28(1): 89-106.

Perry-Smith, J. E. 2006. Social yet creative: The role of social relationships in facilitating individual creativity. Academy of Management Journal, 49(1): 85-101.

Podolny, J. M. & Baron, J. N. 1997. Resources and relationships: Social networks and mobility in the workplace. American Sociological Review, 62(5): 673.

Raymond, E. S. 1999. The cathedral & the bazaar : musings on Linux and open source by an accidental revolutionary (1st ed.). Cambridge, Mass.: O'Reilly.

Reagans, R. & McEvily, B. 2003. Network structure and knowledge transfer: The effects of cohesion and range. Administrative Science Quarterly, 48(2): 240-267.

Scott, J. 2000. Social network analysis : a handbook (2nd ed.). London ; Newbury Park: Sage Publications.

Sia, C. L., Tan, B. C. Y., & Wei, K. K. 2002. Group polarization and computer-mediated communication: Effects of communication cues, social presence, and anonymity. Information Systems Research, 13(1): 70-90.

Sparrowe, R. T., Liden, R. C., Wayne, S. J., & Kraimer, M. L. 2001. Social networks and the performance of individuals and groups. Academy of Management Journal, 44(2): 316-325.

Tapscott, D. & Williams, A. D. 2006. Wikinomics: How Mass Collaboration Changes Everything: Portfolio Hardcover.

Travers, J. & Milgram, S. 1969. An Experimental Study of the Small World Problem. Sociometry, 32: 425-443.

Tsai, W. P. 2001. Knowledge transfer in intraorganizational networks: Effects of network position and absorptive capacity on business unit innovation and performance. Academy of Management Journal, 44(5): 996-1004.

von Hippel, E. & von Krogh, G. 2003. Open source software and the "private-collective" innovation model: Issues for organization science. Organization Science, 14(2): 209-223.

Wagner, C. & Majchrzak, A. 2006. Enabling customer-centricity using wikis and the wiki way. Journal of Management Information Systems, 23(3): 17-43.

Walker, G., Kogut, B., & Shan, W. J. 1997. Social capital, structural holes and the formation of an industry network. Organization Science, 8(2): 109-125.

Wasko, M. M. & Faraj, S. 2000. "It is what one does": why people participate and help others in electronic communities of practice. Journal of Strategic Information Systems, 9(2-3): 155-173.

Wasko, M. M. & Faraj, S. 2005. Why should I share? Examining social capital and knowledge contribution in electronic networks of practice. MIS Quarterly, 29(1): 35-57.

Wasserman, S. & Faust, K. 1994. Social network analysis : methods and applications. Cambridge, [England] ; New York, NY: Cambridge University Press.

Wilkinson, D. M. & Huberman, B. A. 2007. Assessing the value of cooperation in Wikipedia. Palo Alto, CA: HP Labs.

Zhang, X. M. & Zhu, F. 2007. Group Size and Incentive to Contribute: A Natural Experiment at Chinese Wikipedia. Paper presented at the Workshop of Information Systems Economics (WISE), Montreal, Canada.

Figure 1: Two-mode network of articles and editors

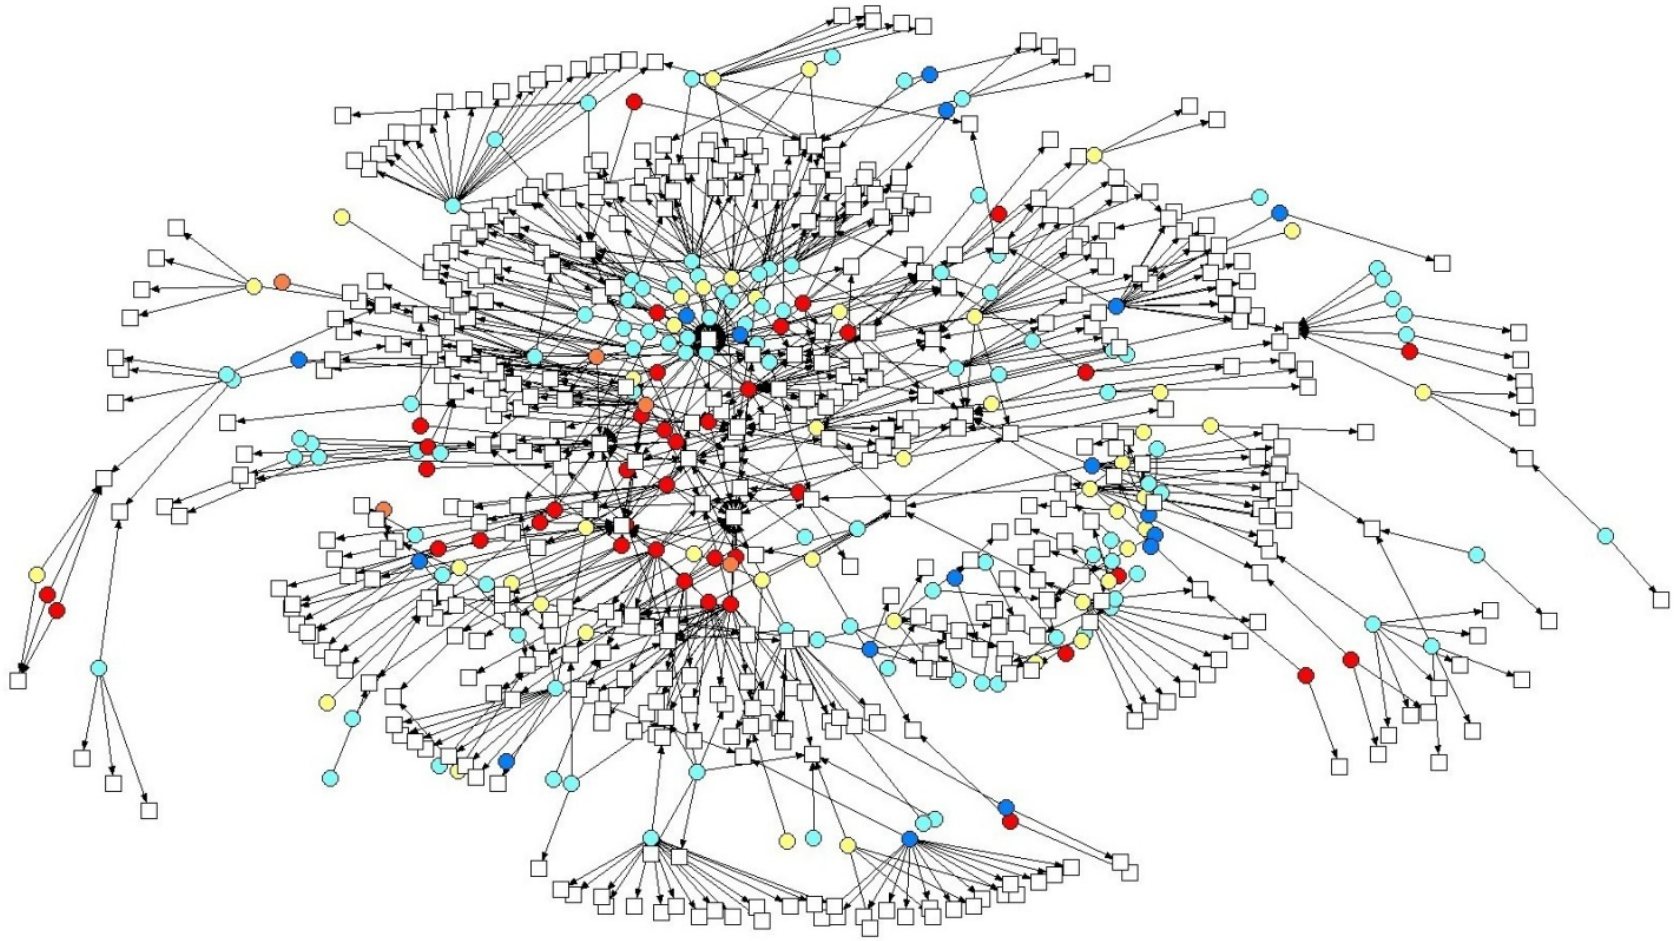

Squares = editors
Circles = articles
Red = Featured Articles
Orange = A-quality Articles
Yellow = Good Articles
Light Blue = B-quality Articles
Dark Blue = Start-quality articles

Table 1: Correlation Matrix

		1	2	3	4	5	6	7	8	9	10	11	12
1	Article Quality	1.000	-.368**	.314**	.317**	.347**	.361**	.230**	-.046	.091	-.084	.457**	.290**
2	Importance	-.368**	1.000	.218**	.152**	.042	-.306**	-.058	.247**	.304**	-.344**	.118*	-.046
3	Traffic	.314**	.218**	1.000	.809**	.730**	-.041	.249**	.387**	.467**	-.438**	.507**	.178**
4	Editors (Art)	.317**	.152**	.809**	1.000	.826**	-.028	.283**	.390**	.534**	-.455**	.479**	.204**
5	Editors (Discussion)	.347**	.042	.730**	.826**	1.000	.138*	.508**	.317**	.343**	-.332**	.372**	.191**
6	Avg Edits/ Editor (Art)	.361**	-.306**	-.041	-.028	.138*	1.000	.480**	-.297**	-.396**	.315**	.117*	.107
7	Avg Edits/ Editor (Discussion)	.230**	-.058	.249**	.283**	.508**	.480**	1.000	.053	.069	-.123*	.202**	.172**
8	Anonymous (Discussion)	-.046	.247**	.387**	.390**	.317**	-.297**	.053	1.000	.614**	-.416**	.137*	.040
9	Anonymous (Art)	.091	.304**	.467**	.534**	.343**	-.396**	.069	.614**	1.000	-.613**	.250**	.109
10	Article Age	-.084	-.344**	-.438**	-.455**	-.332**	.315**	-.123*	-.416**	-.613**	1.000	-.244**	-.090
11	Degree Centrality	.457**	.118*	.507**	.479**	.372**	.117*	.202**	.137*	.250**	-.244**	1.000	.636**
12	Eigenvector Centrality	.290**	-.046	.178**	.204**	.191**	.107	.172**	.040	.109	-.090	.636**	1.000

Table 2: Results								
	MODEL 1				MODEL 2			
	Est.	SE	Wald	Sig.	Est.	SE	Wald	Sig.
Traffic	0.00**	0.00	7.76	0.005	0.00	0.00	2.26	0.13
Google Rank	0.43**	0.15	7.97	0.005	0.43**	0.16	7.51	0.01
Importance	-0.82***	0.12	49.08	0.000	-0.91***	0.12	53.72	0.00
Editors (A)	0.00	0.00	1.17	0.279	0.00	0.00	0.15	0.69
Editors (D)	0.00	0.00	0.08	0.776	0.00	0.00	0.68	0.41
Edits/Editor (A)	0.63***	0.12	28.59	0.000	0.61**	0.12	24.35	0.00
Edits/Editor (D)	-0.06	0.05	1.67	0.196	-0.08	0.05	2.16	0.14
Anon (A)	4.15***	1.04	16.00	0.000	3.47***	1.06	10.72	0.00
Anon (D)	-2.06**	0.75	7.54	0.006	-1.46*	0.76	3.71	0.05
Deg. Cent.					0.00**	0.00	7.29	0.01
Eig. Cent.					0.33***	0.09	12.60	0.00
Pseudo R-Square								
Cox and Snell			0.379					0.474
Nagelkerke			0.407					0.510

Table 3: Descriptive Statistics

	Minimum	Maximum	Mean	Std. Deviation
Traffic	22.00	946121.00	162908.15	169434.29
Google Rank	0.00	1.00	0.60	0.49
Importance	1.00	4.00	2.96	0.72
Editors (A)	6.00	2993.00	540.36	537.35
Editors (D)	2.00	534.00	56.20	71.25
Edits/Editor (A)	0.22	7.41	2.13	0.82
Edits/Editor (D)	1.00	18.68	2.45	1.95
Deg. Cent.	0.00	2796131.00	190368.22	395278.68
Eig. Cent.	0.00	101.64	1.24	8.11
Anon (A)	0.00	0.66	0.48	0.11
Anon (D)	0.00	0.60	0.27	0.13
	Frequency	Percent	Cumulative Percent	
Start-quality	52.00	17.45	17.45	
B-quality	145.00	48.66	66.11	
Good Article	54.00	18.12	84.23	
A-quality	7.00	2.35	86.58	
Featured Article	40.00	13.42	100.00	
Total	298.00	100.00	100.00	

Appendix A: Featured Article Criteria

A featured article exemplifies our very best work and features professional standards of writing and presentation. In addition to meeting the requirements for all Wikipedia articles, it has the following attributes.

1. It is—

(a) well-written: its prose is engaging, even brilliant, and of a professional standard;
(b) comprehensive: it neglects no major facts or details;
(c) factually accurate: claims are verifiable against reliable sources, accurately represent the relevant body of published knowledge, and are supported with specific evidence and external citations; this requires a "References" section in which sources are listed, complemented by inline citations where appropriate;
(d) neutral: it presents views fairly and without bias; and
(e) stable: it is not subject to ongoing edit wars and its content does not change significantly from day to day, except in response to the featured article process.

2. It follows the style guidelines, including the provision of:

(a) a lead—a concise lead section that summarizes the topic and prepares the reader for the detail in the subsequent sections;
(b) appropriate structure—a system of hierarchical headings and a substantial but not overwhelming table of contents (see section help); and
(c) consistent citations—where required by Criterion 1c, consistently formatted inline citations using either footnotes[1] or Harvard referencing (Smith 2007, p. 1) (see citing sources for suggestions on formatting references; for articles with footnotes, the meta:cite format is recommended).

3. Images. It has images and other media where appropriate, with succinct captions and acceptable copyright status. Non-free images or media must satisfy the criteria for inclusion of non-free content and be labeled accordingly.
4. Length. It stays focused on the main topic without going into unnecessary detail (see summary style).