

Published:

Arazy O. , Nov O., Patterson R., and Yeo L., 2011, Information Quality in Wikipedia: The Effects of Group Composition and Task Conflict, *Journal of Management Information Systems (JMIS)*, 27(4), pp. 73-100.

**Information Quality in Wikipedia:
The Effects of Group Composition and Task Conflict**

ABSTRACT

The success of Wikipedia demonstrates that self-organizing production communities can produce high quality information-based products. Research on Wikipedia has proceeded largely atheoretically, focusing on: (i) the diversity in members' knowledge bases as a determinant of Wikipedia's content quality, (ii) task-related conflicts occurred during the collaborative authoring process, and (iii) the different roles members play in Wikipedia. We develop a theoretical model that explains how these three factors interact to determine the quality of Wikipedia articles. The results from the empirical study of 96 Wikipedia articles suggest that (a) diversity should be encouraged as the creative abrasion that is generated when cognitively-diverse members engage in task-related conflict leads to higher quality articles, (b) task-conflict should be managed, as conflict – notwithstanding its contribution to creative abrasion – can negatively affect group output, and (c) groups should maintain a balance of both administrative- and content-oriented members, as both contribute to the collaborative process.

Keywords: community-based production, collaboration, Wikipedia, group composition, cognitive diversity, task conflict, members' orientation, information quality.

Information Quality in Wikipedia:

The Effects of Group Composition and Task Conflict

Introduction

Recent years have seen the emergence of a community-based model for collaborative work, whereby a self-organizing online community creates knowledge-based goods [10,29,52,86]. Prominent examples of community-based collaboration are open source software (OSS) projects (e.g. [49]) and the online encyclopedia Wikipedia [4,14]. Community-based groups are often large, and mostly comprised of volunteers who self-manage work processes. Given the lack of clear monetary incentives for participation and the loose governance, the success of such projects seems counter-intuitive. What is it, then, that determines the success of community-based collaboration projects?

We center our investigation on Wikipedia [88], an online encyclopedia which is experiencing remarkable growth and has become the most notable example of the community-based model [10]¹. To study the performance of editor groups within Wikipedia, we focus on the quality of its main outcome, the information provided in its articles, as the primary indicator of group output. Wikipedia is created using wiki technology, a web-based collaborative authoring tool that is designed for openness, anonymity, and egalitarianism [56,87], and therefore suitable for supporting community-based production of knowledge based goods [88]. Wikipedia operates many different projects, defined as the authoring and editing of a particular encyclopedic article

¹ For a comprehensive list of Wikipedia-related academic papers please refer to Wikipedia's page on the topic http://en.wikipedia.org/wiki/Academic_studies_about_Wikipedia

on a wiki page, where the project group is comprised of the set of volunteers who have contributed to the wiki article. Our definition is in agreement with the inclusive notion of ‘a group’ [57] and with the view of organizations as collections of groups [51].

Prior literature identifies three central aspects in Wikipedia’s collaborative authoring process: (i) the cognitive diversity in members’ knowledge and experiences, (ii) task-related conflict, and (iii) the roles members play. Cognitive diversity of an editor group is believed to contribute to the quality of the produced article, in line with the “the wisdom of the crowd” argument [6,42,80]. Task conflict - disagreements related to the group task - in unavoidable Wikipedia’s peer-based collaboration [43,44,46,83,84], and it has dual effect. On one hand, conflict may impede collaboration and have negative consequences on group output [39,44,46]. Conversely, when cognitively-diverse groups engage in task-related conflicts the ‘creative abrasion’ that is generated forces members to re-examine their assumptions and promotes knowledge [55,63], thus having a positive impact on the performance of Wikipedia editor groups [83,84]. Finally, studies examining members’ roles and group composition in self-organized production communities show that the presence of administrative-oriented members enables the group to diffuse and restrain task conflict [14,79], while the inclusion of content-oriented members draws on their domain-specific expertise [4] and directly impact article quality.

The objective of this paper is to develop and test a limited-scope theoretical model that explains how the three important factors highlighted above - cognitive diversity, task conflict, and group members’ (administrative or content) orientation - interact to determine the quality of articles produced by Wikipedia editor groups. Recent years have produced a large number of studies on Wikipedia; yet, this investigation has proceeded largely atheoretically. Our work is

informed by these many Wikipedia studies, and seeks to ground their observations in theory of group work.

The unique setting of online peer production [52] and the distinctive affordances of the enabling wiki technology [87,84], raise the need to develop refined theories of wiki-based collaboration [60]. Extant theoretical frameworks developed in the context of traditional organizational work are not well suited for explaining Wikipedia work processes. First, some important constructs are only relevant in the context of online self-organizing production groups (e.g. the construct of administrative/content-orientation under investigation here, or ‘cross-thread connectivity’ in Kuk [47]). Second, the relationships between constructs as articulated in extant framework may not be appropriate in the Wikipedia context. Third, no single framework ties together the various constructs we are interested in.

Our empirical method relies on Wikipedia’s system logs. Harvesting these logs can reveal important insights regarding members’ ongoing behavior in its natural setting [39]. While this approach is constrained by the available recorded data, it has advantages over laboratory experiments (investigating the phenomenon in its natural settings) and survey studies (capturing the actual ongoing behavior of all groups’ members) that we believe outweigh its limitations.

The paper continues as follows: the next section reviews related work on cognitive diversity, task conflict, members’ orientation, and information quality; the following section develops hypotheses regarding the relationships between these constructs; we then describe the research methodology and continue to present the results; the section that follows discusses our findings, and highlights implications for theory and practice; and the final section concludes the paper, pointing to future research directions.

Related Studies

Earlier we identified three central themes in Wikipedia's collaborative authoring process. Two themes refer to the composition of Wikipedia editor groups: the inclusion of specific types of members and the diversity in members' knowledge bases. The third theme refers to group work processes: task-related conflict. Here we review prior studies related to these three themes, as well as works related to the group output (the quality of the produced Wikipedia article), focusing primarily on studies of Wikipedia and on works in related areas: virtual teams, online communities, and OSS projects.

The settings for studies in these related areas differ from Wikipedia in some important aspects. Although content in Wikipedia is created by distributed groups, unlike corporate virtual teams, participation in Wikipedia is voluntary and governance is egalitarian [10]. Wikipedia is a production environment where groups work toward a well-defined output and thus Wikipedia differs from many online communities that primarily serve as a place for deliberation. Wikipedia is also somewhat different from other forms of community-based production, such as OSS projects, in its underlying technological infrastructure, namely the affordances of wikis. Wikis are designed for peer-based collaboration [56,87,88] and remove the workflow constraints associated with traditional knowledge management systems [8].

Group Composition

Group composition is an important determinant of project success in community based projects [77] and there is an extensive body of literature that describes the core-periphery structure in online communities (and specifically peer production communities). The core-

periphery structure is defined in terms of members' activity in the project, where a small group of highly active core members are responsible for most of the contribution to the project [19,47,59,65]². Research on the core-periphery structure is descriptive in nature and does not provide insights on the expected effects when communities deviate from this structure. This line of research has focuses on a single dimension of the group's composition - activity - and does not explain how a member's activity, in itself, would determine the quality of the group's output.

There are two important dimensions of group composition that have been investigated [66]: the group's central tendency, i.e. "the proportion of group members who possess a characteristic or the mean level of that characteristic within the group" (p. 367) and the variability within the group in terms of a specific characteristic. We investigate these two dimensions of group composition, focusing on the group's cognitive diversity and on the proportion of members with an administrative (rather than content) orientation.

Group Members Orientation

We center our attention on the focus of group members on a continuum between content- and administrative-orientation. This concept is specific to the context of online communities. Studies of online communities [9,15,45] and Wikipedia [4,11,14,42,71] make the distinction between two types of prototypical users, as described in Table 1 below.

INSERT TABLE 1 HERE

Members of online communities – and specifically, Wikipedia – can volunteer to become administrators. As contributors take on organizational responsibilities, they often develop

² Findings from open source software studies show that, for example, in the Apache project a small core of committed insiders (4% of the members of the project group) was responsible for 88% of the contributions [65] and provided the majority of answers to questions posted on the Apache Usenet help site [48]. Similarly, in Wikipedia core members are responsible for most of the edits [39].

identification with the community and its core values and see themselves as caretakers of this community. They are generally more active than other members, and their activity is often spread across many tasks. The presence of such a core set of administrative-oriented members has become the signature of successful online communities [9,45,93]. A survey of administrators (or ‘owners’) of 284 Internet groups, including both work-related and non-work groups [15], shed light on the roles of these community officials. Administrative-oriented members share responsibilities for both the technology infrastructure and social management tasks, such as controlling for improper use and encouraging appropriate use of the forum. In Wikipedia, administrators are essential for streamlining the production process; they hold decision-making privileges that are associated with their added responsibilities and play an important role in conflict management, quality control, and in defining and enforcing Wikipedia’s norms and procedures [14,79]. While the roles of individual Wikipedia contributors have been investigated extensively [11,14,42,79,84], little is known about administrative-oriented members in article editor groups affect the quality of the produced article.

Another important role in online communities is that of content-oriented contributors who are attracted to the specific topic at the center of discussion, but have less interest in the broader community. Their overall activity is low and sporadic and they usually do not sign on for additional responsibilities. Content-oriented members are an essential component of online communities [9]. Studies of OSS suggest that such occasional members are critical to project success: they contributors provide the variety of ideas that is required to fuel innovation while the administrators select and retain the features that would make up the product [47,52], and similar patterns have been observed in Wikipedia [14,46,79]. In Wikipedia, content-oriented editors are drawn to a particular article because it is closely aligned with their interests and

expertise. The contributions made by these topic-focused members (identified by not registering to Wikipedia) are, on average, of a higher quality than those made by registered members [4].

Cognitive Diversity

Diversity refers to differences among people on “any attribute that another person may use to detect individual differences” [94] (p. 81). Diversity in a group’s composition is often divided to surface-level diversity (differences in demographic characteristics) and deep-level diversity (member differences in educational or functional background, expertise, and knowledge); for a review, see [61,94]. We focus on deep-level diversity, which has the potential to enhance group performance. Group members’ shared mental models and common task knowledge (i.e. low cognitive diversity) have shown to positively affect team success [33,96]. Nonetheless, deep-level diversity can enhance groups’ performance, especially when the task is cognitively complex and requires multiple perspectives or entails creativity [3,32], since cognitive diversity increases the variety of perspectives brought to a problem, creates opportunities for knowledge sharing, and leads to greater creativity [62].

Even though cognitive diversity in online groups could potentially be high, many communities suffer from “cultural tribalism” [41], where people sample a large number of communities and migrate to the ones where they hear what they want to hear, resulting in low cognitive diversity. Thus, cognitive diversity in online communities is only temporary and usually diminishes over time, resulting in dysfunctional communities [41]. To the best of our knowledge, no prior studies of online communities have investigated the effects of cognitive diversity on online work processes and community performance.

Task Conflict

Task conflict refers to differences in viewpoints and opinions pertaining to a group task. There is mixed evidence regarding the effect of task conflict on group performance [18]. For collocated teams, some studies demonstrate the negative effect of task conflict on productivity and satisfaction [75,89], while others [95] show positive effects, especially in tasks that require significant cognitive effort. These positive effects are explained by the benefits derived from differences of opinion about the work being done [22,35], such that the resulting synthesis is superior to the individual perspectives [76]. In virtual collaborations, conflict resolution is of paramount importance because virtual teams are more susceptible to task conflict [34] than collocated teams [31], and thus task conflict may impede the success of virtual collaborations [67]. Empirical evidence from online communities [93] and OSS projects [74,91], suggests that conflict is common. Nonetheless, conflict has been under-investigated in the context of community-based projects.

Conflicts of opinions are inherent to the Wikipedia collaborative authoring process [39,44,46,83]. Decisions about whether a piece of text, link, or image should be included, reordered, or rephrased are all grounds for heated discussion, and resolving these disputes through consensus is the most fundamental discursive work that Wikipedians perform [46]. Even within the bounds of Wikipedia policies, there are often contentions, conflicts, and power plays, where an editor or a group tries to influence the content and claim legitimate control over the article [46]. Conflicts may be observed in the evolution of the article by studying the comments editors append to it [83]. In addition, a ‘talk page’ is associated with each article and provides a place for editors to argue their positions, negotiate disagreements in opinions, and try to reach a consensus [14,46,83]. An analysis of some 4,200 comments in 625 distinct discussion pages revealed that changes appeared and disappeared repeatedly as the community wrestled for

control of content; the most common issues in the discussion pages were task-related conflicts: what content should be added or removed and how information should be structured [39].

Information Quality

Our interest is in the effects of group composition and task conflict on groups' performance, or more specifically on the quality of the produced Wikipedia article. The construct of information quality (IQ) is difficult to pin down. Some definitions highlight the objective view (i.e. information meets the requirements of the particular activity the user is engaged in; e.g. Taylor [81]), while others emphasize a subjective view (i.e. information meets the user's expectations; e.g. Hillgoss and Rieh [30]). For the purposes of this investigation, we adopt the more general definition of quality - "fitness for use" [90]- which encompasses both of these objective and subjective aspects [26].

IQ is a multidimensional construct [81]. There is a large body of research on the concept of IQ and its underlying dimensions [30,53,90,5], and this literature describes a large number of primitive-level attributes of IQ that could be grouped into higher-level categories. In our investigation of Wikipedia article quality, we focus our attention on a restricted set of information quality dimensions, rather than trying to cover the gamut of information quality attributes. For the purpose of this study, we focus on Lee et al.'s [53] IQ categories, employing factors that were frequently used in the studies examined in their survey of the literature: accuracy and objectivity ('intrinsic IQ'), completeness ('contextual IQ'), and representation ('representational IQ').

Theory Development

Given the absence of a theoretical model that explains how group composition variables and task conflict interact to determine content quality there is a need to develop a Wikipedia-centric theoretical model. Recent years have seen numerous studies on Wikipedia, yet much of this investigation has not been grounded in theory. Theories from the related areas of virtual teams and online communities provide some important insights, yet are not directly applicable to our setting. The literature on group processes in virtual teams highlights the essential role of task conflict, but does not consider the composition that characterizes community-based collaboration. Literature on the community-based model, on the other hand, acknowledges the importance of group composition [19] and documents this phenomenon [59], but provides little insight on how this distinctive group composition affects work processes (and specifically task-related conflict) and group output. Moreover, the distinctive affordances of wiki technology promote work processes that differ fundamentally from knowledge creation processes in online communities and virtual teams [60], highlighting the need to develop Wikipedia-centric theory.

In developing our hypotheses we build on two primary literatures. The first is the extensive organizational theory (in both collocated and virtual teams) on the relationships between diversity, conflict, and group performance. The second is the group composition literature in online communities and OSS projects. We explain what effects we expect to persist in Wikipedia and extend these frameworks where necessary.

Cognitive Diversity and Task Conflict Affect Information Quality

Diversity has been shown to have mixed effects on outcome variables [61]. Cognitive diversity has the potential to enhance group performance [62], especially when the task is cognitively complex and requires multiple perspectives or entails creativity [3,32], as is the case

with Wikipedia's collaborative authoring. Conflict has been viewed as a central variable that explains the relationships between diversity and group performance [32,73].

The theory of groups describes the dual effects of task-related conflict on group performance [18]; while conflict can impede collaboration [75,89], it can contribute to group performance [37,95], especially in complex tasks that require significant cognitive effort [22,35]. The explanation for these positive effects is that task conflict (a) makes members more receptive to new information, (b) fosters a deeper understanding of task issues, (c) increases the range of alternatives considered, (d) motivates members to question fundamental assumptions and engage in cognitive restructuring, (e) allows assumptions and recommendations to be evaluated systematically, and (f) eventually leads to novel insights and enhanced problem solving [2,73,97]. For example, Pelled et al. [73] found that task conflict had a positive association with cognitive task performance in work groups and Amason [2] reported that task (or cognitive) conflict is positively related to the quality of group decisions.

When groups are diverse in terms of expertise and viewpoints regarding the task being done, task conflict is most effective [22,35], and the synthesis that emerges from the negotiation of viewpoints is generally superior to the individual perspectives [76]. Recently, a number of studies on organizational learning and innovation have stressed the positive role of creative friction - or abrasion - in facilitating innovation and learning [55,63]. Thus, given a cognitively diverse group of people, task conflict can contribute to the quality of the information-based group output by encouraging members to reassess existing knowledge, challenge others' viewpoints, and come up with creative solutions. In online communities, despite the fact that interaction is computer-mediated, members' postings reveal their knowledge and views; the fact that members have complementary knowledge bases still contributes to the cumulative

knowledge[20,40,64]. Thus, cognitive diversity still plays an important role. Evidence suggests that creative abrasion occurs in online communities, and conflicts were found to help in clarifying issues, eventually resolving disagreements, and reaching (an often a better) commonly-accepted solution [25,93]. We expect Wikipedia groups to experience such a creative abrasion since: (a) these editor groups are often diverse in terms of occupation, geography, age, and breadth and depth of previous experience [39], (b) the collaborative editing task is complex and requires much cognitive effort, and (c) Wikipedia norms and procedures help streamline the collaboration process. Therefore, we propose³:

HYPOTHESIS 1. Task conflict will moderate the relationship between cognitive diversity and information quality, such that when task conflict is high cognitive diversity would have a positive effect, while when task conflict is low cognitive diversity would have a negative effect on information quality.

Task conflict, beyond its contribution to creative abrasion, may impede collaboration and have negative consequences [75,89]. Conflict in technology-mediated collaboration is often viewed as having negative consequences. Virtual teams may suffer from task conflict, as the coordination of conflict resolution is more difficult virtually, and as a result task conflict impedes

³ While extant organizational theory conceptualizes task conflict as a mediator of the relationship between cognitive diversity and group performances [32,73], this conceptualization may not be suitable for our study of Wikipedia. Theories related to diversity and conflict were developed and tested in settings with a clear temporal sequence in the measurement of constructs: group members' knowledge is measured prior to beginning the specific task; conflict is analyzed during the work on the task, and group performance is assessed at completion. In addition, the group boundary is kept constant for the duration of the task. In the context of Wikipedia, however, such a temporal sequence could not easily be established. Wikipedia editor groups work on articles over prolonged time period, gradually developing articles' contents through an ongoing exchange of opinions, and members' knowledge is likely to evolve over this time. Moreover, over time, the group changes, new members join and others leave, and thus the group's cognitive diversity changes even if the knowledge base of each member was to remain constant. Hence, in our study of Wikipedia we model the relation between cognitive diversity and task conflict as a moderation effect: the interaction between group members' diversity in knowledge-base and the task-related conflict experienced when working on the article impacts the quality of the article.

project success [67]. In Wikipedia, groups become preoccupied with minor semantic issues and have difficulty resolving them [39], disagreements often escalate to heated debates [44], and conflict is generally perceived as detrimental [46]. Thus, we expect that – like the reported relations in collocated and virtual teams – the direct effect of task conflict on Wikipedia article quality (beyond its interaction effect with cognitive diversity) will be negative, and we propose:

HYPOTHESIS 2. Task conflict will be negatively related to information quality.

Group Members' Orientation, Task Conflict, and Information Quality

We conceptualize a group members' orientation as a single construct on a continuum between administrative- and content-orientation. The more there are members who act as administrators, the more the group leans toward administrative-orientation; the more there are members who are occasional contributors attracted to the article's topic, the more the group gravitates toward content-orientation.

The inclusion of administrative-oriented members who have gone through the Wikipedia socialization process can enable groups to better manage task conflict, and thus indirectly impact the article's quality. Norms play a critical role in providing a shared terminology and regulating behavior in online communities [9,71,93], and specifically in Wikipedia [79,83]. By relying on policies, editors can more easily interpret complex situations and legitimate their actions [46], as well as work through and resolve content disputes [11,84]. The rules and norms that guide Wikipedia collaboration consist of over one hundred style, presentation, and formation rules, as well as an extensive set of rules for managing the work process (e.g. conflict resolution) [83,84]. Administrators have gone through the Wikipedia socialization processes [11,14,79,83] and in addition play an important role in developing, promoting, and enforcing policies [14,71,79].

Thus, administrative-oriented members share Wikipedia culture and norms more than other members do [14,46], and groups that include high concentration of administrative-oriented editors are more likely to agree on the collaborative authoring process.

In addition, administrators use policies to actively help diffuse conflicts once they emerge [44,46]⁴. As has been observed in collocated teams [36] and OSS projects [54], administrative-oriented Wikipedians play an active role in facilitating conflict resolution [71]. These administrators serve on arbitration and mediation committees, which help diffuse extended conflicts that the community has been unable to resolve [14]. In sum, groups that have a high concentration of administrative-oriented members are likely to experience less conflict, since these members share a common understanding of Wikipedia work processes and can more easily diffuse conflicts and prevent them from escalating. Therefore,

HYPOTHESIS 3. The more group members gravitate towards administrative-orientation, the lower task-conflict the group experiences.

Content-oriented members also play an important role in Wikipedia, and the inclusion of such members can have a direct positive impact on the quality of article the group produces. A central tenet in the user innovation literature[85] is that occasional contributors are essential in innovation processes, since they sometimes possess expertise that surpass the skills of those in official roles. In addition, these occasional contributors are not entrenched in any one opinion, and likely to challenge old conceptions. Findings from OSS projects [47,52] document a similar phenomenon. In the context of online knowledge sharing communities, it has been shown that

⁴ For more detail on Wikipedia's conflict resolution mechanisms please refer to http://en.wikipedia.org/wiki/Wikipedia:Dispute_resolution

content-oriented members who restrict their contributions to specific domains provide higher quality answers than members who contribute across the board. In Wikipedia, content-oriented contributors are drawn to the article's specific topic because of their interest and expertise in the area. They have little interest in Wikipedia beyond the specific article and often choose not to register. A recent study explains the finding that anonymous contributions are of a higher quality in comparison to the contributions of registered members by arguing that these anonymous individuals are often domain experts, while members in official roles are committed to Wikipedia at large, but often lack domain expertise [4]. Thus, we expect a group that includes content-oriented members to produce relatively high quality articles and we hypothesize:

HYPOTHESIS 4. The more group members gravitate towards content-orientation, the higher the quality of the article the group produces.

Controls

In order to control for the effects of exogenous factors, we included a number of control variables. We control for group size, as the number of contributing editors may impact quality in two ways. First, since large groups represent a broader set of opinions and knowledge bases, larger group size may be associated with cognitive diversity and may impact information quality [80]. Second, editors also act as quality assurers; the more editors working on the article, the more rapidly errors and incomplete information will be discovered and corrected (reflecting the principle of 'given enough eyeballs, all bugs are shallow' [74]). Evidence for the relationship between group size and article quality is provided in a recent study that found high correlations between Wikipedia article length and the nomination of articles as high-quality (i.e. 'featured') articles [13]. We also control for factors related to the Wikipedia article. Wikipedia articles

mature slowly, often in small increments. It is through an on-going refinement process that content develops and quality improves. The age of the article might, thus, reflect the article's maturity and may be associated with quality. Similarly, the amount of editorial activity applied to an article could impact the article's quality. Finally, the article length may reflect the completeness of the article and be linked to article quality. We therefore control for these three factors: article age, activity, and length (see [29] for similar controls applied in the study of OSS projects). Lastly, we included in our model control paths from members' orientation to cognitive diversity and from cognitive diversity to information quality.

Research Method

We conducted an empirical study of the English version of Wikipedia. The unit of analysis is an article, the authoring of each article represents a distinct group project, and the group is defined by the set of authors who contributed to the article. Since group output in Wikipedia is an encyclopedic entry about a specific topic, group performance was defined by the quality of the article. According to Kane and Fichman [39], a key limitation of prior research in the area is the narrow focus on easily accessible quantitative data and in order to remedy this, IS research into Wikipedia should employ both quantitative 'shallow' data and deeper analysis of collaboration patterns. We adopt this recommendation and utilize both quantitative data from Wikipedia logs and a detailed content analysis of Wikipedia discussions.

We randomly sampled articles from Wikipedia, with the requirement that the article's length is between 200 and 3500 words. An article's topic may impact conflict levels, since some topics are more controversial than others [43]. In order to represent the various Wikipedia topics, we

used a stratified sampling approach, congruent with Wikipedia’s top-level classification⁵ [43], which we organized into a smaller set of six mutually exclusive and collectively exhaustive classes: (a) culture, art, and religion⁶, (b) math, science, and technology⁷, (c) geography and places, (d) people and self, (e) society⁸, and (f) history and events. The length limit excludes undeveloped (‘stub’) articles, and captures approximately 98% of the non-stub articles⁹. We randomly sampled articles from this population; 17 articles from each category (102 in total), using Wikipedia’s “Random Article” feature. After some preliminary analysis, we found that we had to exclude 6 articles: articles that have not passed the early inception phase (of age less than one month), as well as those with two or less editors (the concepts of group composition carries little meaning in dyads), were discarded; in addition, we removed articles that were discontinued by Wikipedia during the study’s duration. We arrived at a sample of 15-17 articles per category and 96 in total.

In order to measure the quality of Wikipedia articles in our sample, we used the following procedure. In December 2006, we created a copy of each article in our sample, and this version of the articles was analyzed by three senior librarians at a major North American university. The librarians used a variety of sources for judging the quality of Wikipedia articles, including

⁵ For a list of Wikipedia top-level categories, please refer to http://en.wikipedia.org/wiki/List_of_overviews

⁶ Our ‘culture, arts, and religion’ class corresponds to the following Wikipedia categories: ‘Culture and the arts’, ‘Religion and belief systems’, and ‘Philosophy and thinking’.

⁷ Our ‘math, science, and technology’ class corresponds to the following Wikipedia categories: ‘Mathematics and logic’, ‘Natural and physical sciences’, and ‘Technology and applied sciences’.

⁸ Our ‘society’ class corresponds to the following Wikipedia categories: ‘Society and social sciences’ and ‘Health and fitness’.

⁹ The lower limit of 200 words excludes stubs, which represent roughly 30% of all Wikipedia articles. Very long articles are rare since Wikipedia’s guidelines suggest that articles be concise, such that Wikipedia articles contain on average 460 words. Thus, the upper limit of 3500 words on article length excludes only extremely lengthy outliers. For details on Wikipedia article length guidelines please refer to http://en.wikipedia.org/wiki/Article_length. For statistics on article length please refer to [13] and see <http://stats.wikimedia.org/EN/TablesWikipediaEN.htm> and http://en.wikipedia.org/wiki/Wikipedia:Size_comparisons.

reports that were prepared by undergraduate students as part of a class assignment¹⁰. After an initial training session, each of the three librarians independently analyzed the articles in a random order by both performing her own analysis of the topic and by using the students' reports. The librarians were asked to pay attention to the particular set of relevant information quality dimensions described above in the section reviewing studies of information quality: accuracy, completeness, objectivity and representation. Once the librarians independently marked down their perceptions of the articles' quality along these dimensions, we sought to reach a consensus on the overall quality of the article, and a researcher facilitated a discussion of the evaluations, where conflicting views were negotiated and a consensus reached (in line with the Delphi methodology [58]). The consensus quality value for each article (on a 7-point Likert scale) was used as the *information quality* measure.

Group composition variables - cognitive diversity and members' orientation - were measured by 'harvesting' Wikipedia, relying heavily on articles' history logs maintained by Wikipedia and contributors' personal pages. The data was cleaned by excluding all non-human authors (i.e. software bots). Since the contents of Wikipedia changes continually, the estimates for all constructs were based on the article version on which quality assessments were made.

¹⁰ Each student was randomly assigned to analyze three articles. When analyzing an article, students had to identify five alternative information sources that are relevant to the article's topic, and were given clear guidelines on considering the authority of these alternative sources, distinguishing between academic sources, government and non-profit organizations, media institutions, corporate, and personal publications (e.g. a personal web page). Students compared the contents of the Wikipedia article to the alternative sources, and based on this analysis they prepared an extensive report that included: ten pieces of evidence (either validating or invalidating the Wikipedia information), open-ended comments, and ratings of their perceptions of the article's quality. Students were instructed to consider factual accuracy, completeness of information, objectivity and biases, and representation (readability, clarity). To encourage diligent work, students were marked on the sources they employed (credit was linked to the authority of sources) and the gravity of the evidence (detecting a factual error received more credit than detecting a typo). Initially, six students were assigned to each article. We only used reports from students who consented to release their work for this study and utilized only articles with a minimum of three student reports. We later tested to make sure that the number of reports employed did not affect the article's quality score.

We operationalized *group members' orientation* as a reflective latent construct¹¹, where high values indicate administrative-orientation and low values denote content-orientation. We measured this construct using four items, in line with the characteristics described in Table 1. First, we calculated the proportion of group members who hold an official ‘administrator’ status within Wikipedia¹². Second, we calculated the average number of edits the article authors have made across Wikipedia at large, where a high value is associated administrative-orientation. Third, since content-oriented members concentrate on a specific article, but their overall contribution to Wikipedia is low, we calculated the ratio of members’ overall Wikipedia edit activity to their article activity; a high ratio indicates an administrative –orientation and a low value denotes content-orientation. Fourth, since administrative-oriented members spread their activity across many articles and content-oriented contributors are characterized by a concentration of few topics [1], we analyzed – for each member – the dispersion in his Wikipedia activity pattern. We used the entropy measure [3,20], $Entropy = -\sum_{i=1}^N P_i \ln P_i$, where N is the number of states (i.e. different articles the member has contributed to), and P_i is the probability of being in state i (i.e. the percentage of edits in that specific article). Entropy is high when a member contributed evenly across the various articles, and is low when he concentrated his activity on few articles. Entropy is well-suited for estimating the extent to which categories are equally represented, as it measures dispersion irrespective of the location of the mean, unlike variance which measures concentration only around the mean [21]. In the organizational

¹¹ Modeling members’ orientation as a formative construct was not appropriate in this case because of concerns related to content coverage, as well as multicollinearity and network effects [17]

¹² Please note that data on editors is based on how they chose to identify themselves, and data might be imprecise when an administrator chooses to edit articles while not registered or when a user contributes under multiple accounts.

literature, entropy has been employed to measure dispersion of power structure [68], categorical variables (background, gender, and race) [73], and functional assignments [20]. We then calculated the average of group members' entropy, where a high value denotes administrative-orientation and low entropy signifies content-orientation.

Cognitive diversity was estimated based on the uniqueness of group members' experiences within Wikipedia. The uniqueness in group members' knowledge (i.e. the extent to which there is no overlap in the interests and expertise members bring to the group) is a key descriptor of the group's cognitive diversity [96]. Our measure for the cognitive diversity of a Wikipedia article editor group was based on the editors' activity in Wikipedia articles other than the specific article at hand, assuming that an editor's contribution to an article reflects his interests and expertise. For each editor group, we created a 2 dimensional matrix, on one dimension listing all the group members and on the other dimension all the Wikipedia articles that were edited by at least one of the members, where each cell indicated whether the editor was active on that specific article (0 or 1). When there is much overlap in members' Wikipedia activity the matrix would be dense, and when members have a unique set of Wikipedia experiences the matrix is sparse. Thus, our measure of cognitive diversity was based on the sparsity of the matrix (i.e. the percent of cells with zeros [78])¹³.

Task conflict was measured by an analysis of articles' Discussion Pages. Wikipedia associates a discussion page with an article's main page. Editors commonly use the discussion page to discuss, argue, and negotiate their views regarding the information that should be

¹³ Sparsity is often used in the analysis of online social networks to measure how connected the network is (c.f. Wellman [92]). The network sparsity metric is analogous to our measure when we consider a pair of group members that have edited the same Wikipedia page as being connected.

included in an article's main page. Thus, the discussion pages serve as the primary mechanisms for managing conflicts in Wikipedia [44,83,84]. Since the discussion page captures the conflicts, challenges, and decisions in the group involved in creating the article, it is recommended as a source for studying the group's dynamics [39]. Three independent assessors analyzed the discussion page of each article in the sample, and rated the level of task conflict working with the Jehn and Mannix scale [37] adapted to the Wikipedia context (see Appendix A for details). The assessors ranked their perceptions of the levels of task conflict evident in the discussion page text (using a 7-point Likert scale)¹⁴. For each item, we employed the average of the three assessors, and inter-rater agreement was calculated using the intra-class correlation agreement (ICC agreement). Agreement was high for all three conflict perception items (ICC above 0.8 [50]). For testing the moderating effect of *task conflict* on the relation between *cognitive diversity* and *information quality* (hypothesis 1) we created an interaction variable, *Cognitive Diversity x Task Conflict*, by mean-centering indicator items before multiplication.

Control variables were operationalized as follows. *Group Size* was measured based on the number of editors who have contributed to the article. *Article Age* was measured by the number of days since the inception of the article. *Article Activity* was extracted from the article's 'History' tab and calculated as the number of revisions (i.e. edits) the article went through. *Article Length* was measured by the number of words included on the article's wiki page.

¹⁴ In order to corroborate the perceptions measures, we investigated an alternative metric of task conflict. The assessors were instructed to analyze the discussion pages and highlight the words related to task conflict. We then counted the number of conflict-related words and calculated their percentage from the overall number of words in the discussion page. After verifying inter-rater agreement, we calculated assessors' average word-count percentage. We found that the word-count metric highly correlates (0.73-0.78) with the perception measures. Furthermore, the pattern of effects remains unchanged when the word-count metric replaces the conflict perception measures, thus further validating our analysis.

Results

Descriptive Statistics

Descriptive statistics are outlined in Table 2 below. The average group included 49 active content contributors, where 15% of the contributors are administrators, 52% regular members, and 33% who have contributed anonymously. Members made on average over 6,200 edits to Wikipedia articles. Cognitive diversity was on average 0.92, reflecting the little overlap in members' activity outside the group. The average article was over 2 years old, went through roughly 90 edits, and included close to 1000 words. Discussion pages' average level of conflict was relatively low, and its variance was quite high. Finally, the quality of articles is moderately high: 4.4 out of 7.

INSERT TABLE 2 HERE

Path Model

Data analysis was conducted using the *Partial Least Squares* (PLS) algorithm [38], which estimates multi-stage path models using composite variables from a number of indicator items. In this respect, the variance-based PLS path modeling is similar to covariance-based structural equation modeling (SEM), because both algorithms estimate complex relations between several latent variables simultaneously. Compared to covariance-based SEM, the PLS can be used with smaller samples, thus it is well-suited for this study. Furthermore, PLS requires fewer assumptions about data distributions, and is robust in case these assumptions are violated [16], making it suitable for dealing with variables of non-normal distribution (some of our variables exhibit power-law distribution).

The ordinal indicator items (*task conflict* and *information quality* perception items) were treated as interval scale variables, assuming that the underlying scale is continuous, but because of the lack of a sophisticated instrument, they could not be measured on an interval scale. The psychometric properties of the instrument were analyzed before examining the data for hypothesis testing. The estimate for composite reliability of the multi-item *task conflict* and *group members' orientation* constructs were 0.99 and 0.94 respectively, well above the recommended threshold of 0.8 [70], thus demonstrating good internal consistency. Convergent validity for all constructs was good. The average variance extracted (AVE) for *task conflict* and *group members' orientation* were 0.98 and 0.79 respectively, substantially greater than the suggested minimum of 0.5 [24], and item loadings were 0.80-0.99, greater than the suggested level of 0.7.

Discriminant validity was examined by comparing the square root of the AVE (RAVE) of a particular construct (presented in Table 3 on the diagonal, in bold) and the correlation between that construct and other latent constructs [23] (presented in the off-diagonal positions of the table), as well as based on item loadings. We found discriminant validity to be good. The constructs' RAVE are 0.89 or higher, RAVE for every construct is higher than the correlation between that construct and all other constructs. The correlations between latent constructs did not exceed the recommended threshold of 0.5, with two exceptions related to the control variables. The correlation between two of the control variables - *group size* and *article activity* - was extremely high (0.98). Due to multi-linearity concerns we were not able to include both controls in the same model, and tested the model twice, alternating between *group size* and *article activity*. We also observed relatively high correlation between *task conflict* and these two control variables (0.63 and 0.67); however, we don't perceive this as a risk to discriminant validity since

the metrics for *group size* and *article activity* directly measure these constructs and *task conflict* was estimated from an alternative source. Item loadings on their corresponding constructs were all above 0.7, and those loadings were substantially higher than any cross-loadings on any other construct, thus further supporting discriminant and convergent validity.

INSERT TABLE 3 HERE

Next, we tested the research model by specifying paths in the PLS structural model corresponding to the model's hypotheses. The significance of structural path estimates was computed using the bootstrapping re-sampling method (with 200 re-samples; cf. Tenenhaus et al.[82]), and the structural model was evaluated based on both the R^2 for each composite variable and the statistical significance of structural paths. All the independent variables were standardized. Figure 1 shows the results of the PLS analysis.

INSERT FIGURE 1 HERE

As shown in Figure 1, a substantial amount of variance in *information quality* was explained by our model ($R^2 = 32\%$). The effects of all paths but one were statistically significant. The results of the PLS analysis demonstrate that the interaction between *cognitive diversity* and *task conflict* was statistically significant (H1)¹⁵; *task conflict* had a direct negative impact on *information quality* (H2); the presence of *group member's orientation* has direct negative effects on both *task conflict* (H3) and *information quality* (H4). The only control variable to have a significant (positive) effect on the outcome variable was *article length*. The only control path to

¹⁵ In order to address the concern raised by Goodhue et al. [28] regarding the statistical power of interaction effects in PLS (as compared to using multiple regressions), we ran a regression analysis and found that the interaction effect (as well as all other hypothesized relations) were still significant in the regression model.

have a significant effect was from *members' orientation to cognitive diversity*. The results when using *article activity* as a control instead of *group size* were indistinguishable¹⁶.

Discussion

The success of community-based projects has attracted significant attention in the research community [86], and Wikipedia has been investigated extensively in recent years [39,88]. Yet, many of the studies on Wikipedia success, and specifically its article quality, have been rather anecdotal [27,79] and our understanding of the factors driving Wikipedia content quality is still limited. In this study we explored the relationships between groups' cognitive diversity, the conflict that arises during the collaborative authoring process, group members' orientation, and the quality of Wikipedia articles created by these groups.

The composition of Wikipedia has been described in previous studies, demonstrating that articles' editor groups are characterized by a composition of few highly active members and many occasional editors that make few contributions [39], similar to the core-periphery structure that characterizes online communities and OSS projects [19,47,59,65]. However, it is not clear how the core-periphery composition affects collaborative work processes and group performance. We argue that the main limitation of research on group composition in community-based projects is the under-specification of the dimensions on which composition is analyzed. To address this gap, we investigate two dimensions of Wikipedia groups' composition: group members' orientation and diversity in members' knowledge¹⁷. We contribute to the literature of

¹⁶ It is worth noting that when running the model without the control variables, all hypotheses were supported and all paths were statistically significant, including H3 related to the un-moderated effect of cognitive diversity on information quality.

¹⁷ Note that the activity-based core-periphery structure is fundamentally different from the composition dimensions we are interested in. Our group composition constructs are defined at the level of the institution (i.e. Wikipedia):

members' roles by formulating the constructs of members (administrative/content) orientation and by proposing measures for this construct. Our model demonstrates three ways in which these group composition constructs, through their interaction with task conflict, impact the quality of the article produced by the group. First, the presence of administrative-oriented members enables the group to diffuse and restrain task conflict [11,14,79,84]. Second, the inclusion of content-oriented members draws on their domain-specific expertise [4] and directly impact article quality. Third, the interaction between cognitive diversity and task conflict generates a creative abrasion [55] that positively affects information quality.

An important contribution of this study is in documenting the effect of *cognitive diversity* – in interaction with *task conflict* – on *information quality*. Prior studies show that diversity of expertise and knowledge has the potential to enhance group performance [62], especially when the task is cognitively complex and requires multiple perspectives or entails creativity[3,32], as is the case with Wikipedia's collaborative authoring. Moreover, in cases where groups are cognitively diverse, conflicts related to the task at hand can produce a creative abrasion that results in more innovative solutions [55,63]. The results of our study show that task conflict moderates the effect of diversity (effect size = 0.48; $p < 0.05$), such that when task conflict is low cognitive diversity is detrimental, but when conflict is high cognitive diversity has a positive effect on information quality. Our findings support earlier results of online groups [25,93], and illustrate the value of creative abrasion in Wikipedia's collaborative authoring. This moderation effect is illustrated in Figure 2 below.

group members' orientation and cognitive diversity were both based on editors' activity patterns across Wikipedia. The core-periphery structure, on the other hand, is defined at the project group level (please refer to [7] for a discussion of the difference between project-level and organization-level core-periphery composition in Wikipedia).

INSERT FIGURE 2 HERE

Task conflict, beyond its positive moderating effect, had a direct negative effect on information quality (effect size = 0.42; $p < 0.01$). This main effect, in the presence of an interaction term, is interpreted as the effect of a given exogenous predictor when the moderator it interacts with is at its mean value [12]. Our findings can inform the organizational literature on group conflict, and offer a possible explanation for the reported mixed effects of task conflict [18]. Note that while conflict has been discussed extensively in previous Wikipedia studies [14,36,44,54,71], to the best of our knowledge this is the first study to demonstrate the relationships between cognitive diversity, task conflict, and article quality.

We also contribute to the growing body of literature on Wikipedia – and more broadly to the online communities’ literature – by demonstrating how the presence of administrative-oriented members contributes to the quality of the information good the group produces. While there is an extensive literature describing the roles of administrators in Wikipedia [44,46,71] and other communities [15], and their contribution in terms of quality assurance, conflict resolution and the development of procedures, there is little evidence to show how these efforts impact product quality. This study contributes to the field by demonstrating how the inclusion of administrative-oriented members helps to restrain task conflict (effect size = 0.34; $p < 0.001$). It is interesting to note that this effect acts as a double-edged sword: on one hand it limits the negative direct effect of task conflict on information quality, and on the other hand it limits the potential for creative abrasion, thus restricting the potential value of cognitive diversity. This contrasts with the all-positive view of administrators in extant online communities literature [14,15,44,46,54,71,79].

A third contribution of this study is in demonstrating the direct positive effect of group’s content-orientation on the quality of Wikipedia articles (effect size was 0.29; $p < 0.05$). Although

this construct has not been investigated to date in the context of Wikipedia, results from the related areas of user innovation [85], OSS projects [47,52], and online communities [1] show that occasional contributors are essential in the design process, since they sometimes possess expertise that surpass the skills of those in official roles. In Wikipedia, Anthony et al. [4] demonstrated that the anonymous contributions (likely by those who are drawn to the article's specific topic) are of a higher quality in comparison to the contributions of registered members. To the best of our knowledge, ours is the first study to demonstrate this effect at the group level. Our findings regarding the direct effects of members' orientation are of special importance, given that most extant research on online communities, and in particular production communities, assumes that greater engagement leads to more successful projects, and thus focuses on the motivational factors driving participation and engagement [69,86,72]. Our findings reveal that this line of research provides only a partial explanation of the factors driving the success of online communities, and that the inclusion of occasional, content-oriented contributors (who are often less engaged and active) is essential to producing quality outputs. Another distinctive aspect of this study is that it investigates performance at the group level; research on the motivations for participation in online communities focuses on individual performance, and the factors contributing to group outputs have been under investigated. This study, thus, helps to fill this gap and explain how members' orientation impacts group output.

Conclusion

This study contributes to the growing research on Wikipedia by developing and testing a framework for the determinants of the quality of group output. The study's primary contribution is in demonstrating the multiple effects of group composition. The presence of occasional editors contributes to the group, as these content-oriented contributors often possess domain-specific

expertise. Administrative-oriented members, on the other hand, can help the group restrain and diffuse task conflict. Lastly, groups that contain members with diverse knowledge bases and experiences may benefit from the variety of opinions, especially in the presence of task conflict that promotes creative abrasion. Our contribution extends beyond the specific Wikipedia context: our formulation of the ‘members’ orientation’ construct and the findings regarding its effects can inform research on online communities, and the formulation of the relationship between cognitive diversity and task conflict contributes to the research on group work.

In addition, our findings have two important implications for practice. First, our findings illustrate that it is difficult to compose groups for effective performance, since opposing ends along the group members’ orientation continuum – administrative- and content-orientation – both play an important role, and it would be challenging to optimize the effects of these opposing forces. Organizers of online production communities should strive to reduce entry barriers in order to entice the participation of occasional content-oriented contributors (which may possess relevant expertise) and increase group diversity; at the same time, they should strive to include some administrative-oriented members that would streamline the production process and help diffuse conflicts. A second important practical implication is for managing task conflict and providing conflict resolution tools. Communities should encourage task-related conflicts, while providing the tools for preventing conflicts from escalating and the mechanisms to resolve conflicts. Some have suggested that the success of Wikipedia can be attributed to the diversity of opinions and knowledge-bases it attracts (i.e. ‘The Wisdom of the Crowds’ argument); our findings suggest that Wikipedia’s extensive conflict resolution mechanisms are responsible for unleashing the potential of this cognitive diversity. Another important aspect of Wikipedia’s conflict management is in preventing task-related conflicts from escalating into destructive

personal conflicts; the Wikipedia community has developed an extensive set of behavior norms for preventing such conflict intensifications. Finally, our data-driven proxies of group composition variables offer a practical way for practitioners to analyze the effects of group composition and determine optimal arrangements for their specific settings.

Our study provides preliminary results and further research is warranted. The study concerns a large number of projects in one setting (i.e. Wikipedia). As such, it provides a useful context for the preliminary validation of our research model, as the setting controls for exogenous factors that might have interfered in an inter-organizational study (e.g. differences in organizational culture and norms). We expect that our findings will generalize to other forms of community-based knowledge production projects, beyond Wikipedia. However, it is possible that specific features of Wikipedia (both in terms of the social mechanisms and underlying technological infrastructure) play a role in allowing the effects of group composition and task conflict (e.g., low entry barriers in Wikipedia may affect group composition). Hence, further research is warranted in order to test whether our findings hold in other types of online communities. Another way in which our study could be extended is in the operationalization and measurement of the outcome variable. We have used university librarians as assessors; however, future studies may help validate our proposed model by using alternative methods for estimating quality. In addition, we have conceptualized information quality as a single construct; we suggest that future research explore specific information quality dimensions.

Finally, our data-driven research method constrained the type of constructs that could be investigated. The availability of comprehensive records of wiki-based collaboration allows us to analyze group members' actual ongoing behavior in its natural settings [39]. Still, the main limitations of this approach are that the measurement of constructs is restricted to available

proxies and that other relevant factors which have not been recorded might be omitted. Alternatively, we could create groups in a laboratory experiment, and test cause and effect relationships between relevant variables; however, it would be difficult to generalize these results to institutions such as Wikipedia. Thus, the proposed data-driven approach has advantages that – we believe – outweigh its limitations, and this approach has been valuable in the study of online communities [47]. That being said, we hope that future research would conduct more comprehensive investigation of the factors explaining the success of community-based projects, include additional relevant variables (for example, members' personal characteristics, e.g. expertise and motivations), and use more direct measures of these constructs.

In conclusion, we believe that Wikipedia can serve as a test bed for studying decentralized IT-enabled production of knowledge-based goods. Research on Wikipedia is in early phases, and the success of Wikipedia-like projects remains to a large extent unexplained. This study extends current understanding of online production communities by enhancing our knowledge of the role of relatively unexplored, yet important, factors – cognitive diversity, task conflict, and group members' orientation. We hope that this study will open the door to future research in the area.

Acknowledgements

We thank Kevin O'Kell, Kathy West, Janet Williamson, and Maureen Beristain for their work on this project. The authors also thank the JMIS anonymous reviewers for their constructive comments on earlier drafts of this paper. This research was funded in part by the Canadian Social Sciences and Humanities Research Council (SSHRC).

Bibliography

- 1 Adamic, L. A.; Zhang, J.; Bakshy, E.; and Ackerman, M. S. Knowledge sharing and Yahoo Answers: Everyone knows something. In *Proceedings of the 17th International Conference*

- on *World Wide Web* Beijing: ACM, 2008, pp. 665-674.
- 2 Amason, A. C. Distinguishing the effects of functional and dysfunctional conflict on strategic decision making: Resolving a paradox for top management teams. *Academy of Management Journal*, 39, 1 (1996), 123-148.
 - 3 Ancona, D. and Caldwell, D. Demography and design: Predictors of new product team performance. *Organization Science*, 3, 3 (1992), 321-341.
 - 4 Anthony, D.; Smith, S. W.; and Williamson, T. Reputation and reliability in collective goods. *Rationality and Society*, 21, 3 (2009), 283-306.
 - 5 Arazy, O. and Kopak, R. On the measurability of information quality. *Journal of the American Society for Information Science and Technology* (2010), forthcoming.
 - 6 Arazy, O.; Morgan, W.; and Patterson, R. Wisdom of the crowds: Decentralized knowledge construction in Wikipedia. In *Proceedings of the 16th Workshop on Information Technologies and Systems* Milwaukee: SSRN, 2006.
 - 7 Arazy, O. and Nov, O. Determinants of Wikipedia quality: The roles of global and local contribution inequality. In *Proceedings of Computer Supported Cooperative Work (CSCW) 2010* Savannah: ACM, 2010, pp. 233-236.
 - 8 Arazy, O.; Stroulia, E.; Ruecker, S.; Arias, C.; Fiorentino, C.; Ganev, V.; and Yau, T. Recognizing contributions in wikis: Authorship categories, algorithms, and visualizations. *Journal of the American Society for Information Science and Technology*, 61, 6 (2010), 1166-1179.
 - 9 Bateman, P. J.; Gray, P. H.; and Butler, B. Community commitment: How affect, obligation, and necessity drive online behaviors. In *Proceedings of the 27th International Conference on Information Systems (ICIS)* Milwaukee: AIS, 2006, pp. 983-1000.
 - 10 Benkler, Y. *The wealth of networks: How social production transforms markets and freedom*. Yale University Press, 2006.
 - 11 Beschastnikh, I.; Kriplean, T.; and McDonald, D. W. Wikipedian self-governance in action: Motivating the policy lens. In *Proceedings of the 2nd International Conference on Weblogs and Social Media* Seattle: AAAI, 2008, pp. 27-35.
 - 12 Bharadwaj, S.; Bharadwaj, A.; and Bendoly, E. The performance effects of complementarities between information systems, marketing, manufacturing, and supply chain processes. *Information Systems Research*, 18, 4 (2007), 437-453.
 - 13 Blumenstock, J. E. Size matters: Word count as a measure of quality on Wikipedia. In *Proceedings of 17th International Conference on World Wide Web* Beijing: ACM, 2008, pp. 1095-1096.
 - 14 Bryant, S.; Forte, A.; and Bruckman, A. Becoming Wikipedian: Transformation of participation in a collaborative online encyclopedia. In *Proceedings of the 2005 International ACM SIGGROUP Conference on Supporting Group Work* Florida: ACM, 2005, pp. 1-10.
 - 15 Butler, B.; Sproull, L.; Kiesler, S.; and Kraut, R. Community effort in online groups: Who does the work and why? In S. Weisband (ed.), *Leadership at a Distance: Research in Technology-Supported Work*. Mahwah: Lawrence Erlbaum Associates, 2007.
 - 16 Cassel, C.; Hackl, P.; and Westlund, A. H. Robustness of partial least-squares method for estimating latent variable quality structures. *Journal of Applied Statistics*, 26 (1999), 435-446.

- 17 Cenfetelli, R. and Bassellier, G. Interpretation of formative measurement in IS research. *MIS Quarterly*, 33, 4 (2009), 689-707.
- 18 Chidambaram, L. and Carte, T. Diversity: Is there more than meets the eye? A longitudinal study of the impact of technology support on teams with differing diversity. In *Proceedings of the 38th Hawaii International Conference on System Sciences* Big Island: IEEE Computer Society, 2005.
- 19 Crowston, K.; Wei, K.; Li, Q.; and Howison, J. Core and periphery in free/libre and open source software team communications. In *Proceedings of the 39th Annual Hawaii International Conference on System Sciences* Kauai: IEEE Computer Society, 2006.
- 20 Cummings, J. N. Work group, structural diversity, and knowledge sharing in a global organization. *Management Science*, 50, 3 (2004), 352-364.
- 21 Ebrahimi, N.; Maasoumi, E.; and Soofi, E. Comparison of entropy and variance orderings. *Journal of Econometrics*, 90 (1999), 317-336.
- 22 Eisenhardt, K. and Schoonhoven, C. Organizational growth: Linking founding team, strategy, environment, and growth among U.S. semiconductor ventures, 1978-1988. *Administrative Science Quarterly*, 35 (1990), 504-529.
- 23 Fornell, C. and Larcker, D. A second generation of multivariate analysis: Classification of methods and implications for marketing research. In M. J. Houston (ed.), *Review of Marketing*. Chicago: American Marketing Association, 1987.
- 24 Fornell, C. and Larcker, D. F. Evaluating structural equation models with unobservable variables and measurement error. *Journal of Marketing Research*, 15 (1981), 39-50.
- 25 Franco, V.; Piirto, R.; Hu, H. Y.; Lewenstein, B. V.; Underwood, R.; and Vidal, N. K. Anatomy of a flame: Conflict and community building on the Internet. *Technology and Society Magazine*, 14 (1995), 12-21.
- 26 Ge, M. and Helfert, M. A Review of Information Quality Research - Develop a Research Agenda. In *Proceedings of the 12th International Conference on Information Quality* Cambridge:, 2007.
- 27 Giles, J. Internet encyclopedias go head to head. *Nature*, 438, 15 (2005), 900-901.
- 28 Goodhue, D.; Lewis, W.; and Thompson, R. Statistical power in analyzing interaction effects: Questioning the advantage of PLS with product indicators. *Information Systems Research*, 18, 2 (2007), 211-227.
- 29 Grewal, R.; Lilien, G. L.; and Mallapragada, G. Location, location, location: How network embeddedness affects project success in open source systems. *Management Science*, 52 (2006), 1043-1056.
- 30 Hilligoss, B. and Rieh, S. Y. Developing a unifying framework of credibility assessment: Construct, heuristics, and interaction in context. *Information Processing & Management*, 44, 4 (2008), 1467-1484.
- 31 Hinds, P. J. and Mortensen, M. Understanding conflict in geographically distributed teams: The moderating effects of shared identity, shared context, and spontaneous communication. *Organization Science*, 16, 3 (2005), 290-307.
- 32 Hoffman, L. and Maier, N. Quality and acceptance of problem solutions by members of homogeneous and heterogeneous groups. *Journal of Abnormal and Social Psychology*, 62

- (1961), 401-407.
- 33 Hoopes, D. G. and Postrel, S. Shared knowledge, 'glitches,' and product development performance. *Strategic Management Journal*, 20 (1999), 837-865.
 - 34 Jarvenpaa, S.; Knoll, K.; and Leidner, D. Is there anybody out there? Antecedents of trust in global virtual teams. *Journal of Management Information Systems*, 14, 4 (1998), 29-64.
 - 35 Jehn, K. A. A multimethod examination of the benefits and detriments of intragroup conflict. *Administrative Science Quarterly*, 40 (1995), 256-282.
 - 36 Jehn, K. A qualitative analysis of conflict types and dimensions in organizational groups. *Administrative Science Quarterly*, 42, 3 (1997), 530-557.
 - 37 Jehn, K. and Mannix, E. The dynamic nature of conflict: A longitudinal study of intragroup conflict and group performance. *Academy of Management Journal*, 44 (2001), 236-251.
 - 38 Jöreskog, K. G. and Wold, H. The ML and PLS techniques for modeling with latent variables: Historical and comparative aspects. In K. G. Jöreskog and H. Wold (eds.), *Systems under indirect observation: Causality, structure, prediction. Part I*. Amsterdam:North-Holland, 1982, pp. 263-270.
 - 39 Kane, G. C. and Fichman, R. G. The shoemaker's children: Using wikis to improve IS research, teaching, and publication. *MIS Quarterly*, 33, 1 (2009), 1-22.
 - 40 Kiesler, S. and Cummings, J. What do we know about proximity and distance in work groups? In P. Hinds and S. Kiesler (eds.), *Distributed Work*. Cambridge: MIT Press, 2002, pp. 57-82.
 - 41 Kitchin, R. *Cyberspace: The world in the wires*. Chichester: Wiley, 1998.
 - 42 Kittur, A.; Chi, E. H.; Pendelton, B. A.; Suh, B.; and Mytkowicz, T. Power of the few vs. wisdom of the crowd: Wikipedia and the rise of the bourgeoisie. In *Proceedings of alt.CHI*, San Jose: ACM, 2007.
 - 43 Kittur, A.; Suh, B.; and Chi, E. H. What's in Wikipedia? Mapping topics and conflict using socially annotated category structure. In *27th Annual CHI Conference on Human Factors in Computing Systems* Boston: ACM, 2009, pp. 1509-1512.
 - 44 Kittur, A.; Suh, B.; Pendleton, B. A.; and Chi, E. H. He says, she says: conflict and coordination in Wikipedia. In *Proceedings of the 25th Annual ACM Conference on Human Factors in Computing Systems* San Jose: ACM, 2007, pp. 453-462.
 - 45 Kollock, P. and Smith, M. Managing the Virtual Commons: Cooperation and conflict in computer communities. In S. Herring (ed.), *Computer-Mediated Communication: Linguistic, Social, and Cross-Cultural Perspectives*. Amsterdam: John Benjamins, 1996 pp. 109-128.
 - 46 Kriplean, T.; Beschastnikh, I.; McDonald, D. W.; and Golder, S. A. Community, consensus, coercion, control: cs*w or how policy mediates mass participation. In *Proceedings of the 2007 International Conference on Supporting Group Work* Sanibel Island: ACM, 2007, pp. 167-176.
 - 47 Kuk, G. Strategic interaction and knowledge sharing in the KDE developer mailing list. *Management Science*, 52 (2006), 1031-1042.
 - 48 Lakhani, K. R. and von Hippel, E. How open source software works: "free" user-to-user assistance. *Research Policy*, 32, 6 (2003), 923-943.

- 49 Lakhani, K. R. and Wolf, R. G. Understanding motivation effort in free/open source software projects. In J. Feller, et al. (eds.), *Perspectives on Free and Open Source Software*. Cambridge: MIT Press, 2005, pp. 3-22.
- 50 Landis, J. R. and Koch, G. G. The measurement of observer agreement for categorical data. *Biometrics*, 33, 1 (1977), 159-174.
- 51 Lawler, E. J. Affective attachments to nested groups: A choice process theory. *American Sociological Review*, 57 (1992), 327-339.
- 52 Lee, G. and Cole, R. From a firm-based to a community-based model of knowledge creation. *Organization Science*, 14, 6 (2003), 633-649.
- 53 Lee, Y. W.; Strong, D. M.; Kahn, B. K.; and Wang, R. Y. AIMQ: A methodology for information quality assessment. *Information & Management*, 40, 2 (2002), 133-146.
- 54 Lehmann, F. FLOSS developers as a social formation. *First Monday*, 9, 11 (November 2004).
- 55 Leonard-Barton, D. *Wellsprings of Knowledge: Building and Sustaining the Sources of Innovation*. Cambridge: Harvard Business Press, 1995.
- 56 Leuf, B. and Cunningham, W. *The Wiki Way: quick collaboration on the web*. Boston: Addison-Wesley, 2001.
- 57 Levine, J. M. and Moreland, R. L. Small groups. In D. Gilbert et al. (eds.), *The handbook of social psychology*. Boston: McGraw-Hill, 1998, pp. 22-24.
- 58 Linstone, H. A. and Turoff, M. *The Delphi Method: Techniques and Applications*. London: Addison-Wesley, 1975.
- 59 Long, Y. and Siau, K. Network structures in open source software development teams. *Journal of Database Management*, 18, 2 (2007), 25-40.
- 60 Majchrzak, A. Comment: Where is the theory in wikis? *MIS Quarterly*, 33, 1 (2009), 18-20.
- 61 Mannix, E. and Neale, M. What differences make a difference? *Psychological Science in the Public Interest*, 6, 2 (2005), 31-55.
- 62 March, S. Exploration and exploitation in organizational learning. *Organization Science*, 2 (1991), 71-87.
- 63 Matsuo, M. Customer orientation, conflict, and innovativeness in Japanese sales departments. *Journal of Business Research*, 59 (2006), 242-250.
- 64 Mitchell, R. and Nicholas, S. Knowledge creation in groups: The value of cognitive diversity, transactive memory and open-mindedness norms. *Electronic Journal of Knowledge Management*, 4, 1 (2006), 67-74.
- 65 Mockus, A.; Fielding, R. T.; and Herbsleb, J. D. Two case studies of open source software development: Apache and Mozilla. *ACM Transactions on Software Engineering and Methodology*, 11 (2002), 309-346.
- 66 Moreland, R. L. and Levine, J. M. Group composition: Explaining similarities and differences among group members. In M. A. Hogg and J. Cooper (eds.), *The SAGE Handbook of Social Psychology*. London: SAGE Publications, Ltd, 2003.
- 67 Mortensen, M. and Hinds, P. J. Conflict and shared identity in geographically distributed teams. *International J. Conflict Management*, 12, 3 (2001), 212-238.
- 68 Murphy, D. C. and Hasenjaeger, J. T. Entropy as a measure of decentralization. In

- Proceedings of the American Institute for Decision Sciences* Boston., 1973, pp. 67-69.
- 69 Nov, O. What motivates Wikipedians. *Communications of the ACM*, 50, 11 (2007), 60-64.
- 70 Nunnally, J. C. *Psychometric Theory*. New York: McGraw Hill, 1978.
- 71 O'Neil, M. *Cyberchiefs: Autonomy and Authority in Online Tribes*. London: Pluto Press, 2009.
- 72 Oreg, S. and Nov, O. Exploring motivations for contributing to open source initiatives: The roles of contribution context and personal values. *Computers in Human Behavior*, 24, 5 (2008), 2055-2073.
- 73 Pelled, L.; Eisenhardt, K.; and Xin, K. Exploring the black box: An analysis of workgroup diversity, conflict, and performance. *Administrative Science Quarterly*, 44 (1999), 1-28.
- 74 Raymond, E. *The Cathedral and the Bazaar*. O'Reilly, Sebastopol, CA, 2001.
- 75 Saavedra, R.; Earley, P. C.; and Van Dyne, L. Complex interdependence in task-performing groups. *Journal of Applied Psychology*, 78 (1993), 61-72.
- 76 Schwenk, C. Conflict in organizational decision making: An exploratory study of its effects in for-profit and not-for-profit organizations. *Management Science*, 36 (1990), 436-448.
- 77 Singh, P. V.; Tan, Y.; and Mookerjee, V. Social capital, structural holes and team composition: Collaborative networks of the open source community. In *Proceedings of the International Conference of Information Systems* Montreal: AIS, 2007.
- 78 Stoer, J. and Bulirsch, R. *Introduction to Numerical Analysis*. Berlin: Springer-Verlag, 2002.
- 79 Stvilia, B.; Twidale, M. B.; Smith, L. C.; and Gasser, L. Information quality work organization in Wikipedia. *Journal of the American Society for Information Science and Technology*, 59, 6 (2008), 983-1001.
- 80 Surowiecki, J. *The Wisdom of Crowds*. New York: Random House, 2005.
- 81 Taylor, R. S. *Value-added processes in information systems*. Norwood: Ablex Publishing, 1986.
- 82 Tenenhaus, M.; Vinzi, V. E.; Chatelin, Y. M.; and Lauro, C. PLS path modeling. *Computational Statistics and Data Analysis*, 48 (2005), 159-205.
- 83 Viegas, F.; Wattenberg, M.; and Dave, K. Studying cooperation and conflict between authors with history flow visualizations. In *Proceedings of the SIGCHI conference on Human factors in computing systems* Vienna: ACM, 2004, pp. 575-582.
- 84 Viegas, F. B.; Wattenberg, M.; Kriss, J.; and van Ham, F. Talk before you type: Coordination in Wikipedia. In *Proceedings of the 40th annual Hawaii International Conference on System Sciences*, Waikoloa: IEEE Computer Society, 2007, pp. 78.
- 85 von Hippel, E. *Democratizing innovation*. Cambridge: MIT Press, 2005.
- 86 von Krogh, G. and von Hippel, E. The promise of research on open source software. *Management Science*, 52, 7 (2006), 975-983.
- 87 Wagner, C. Wiki: A technology for conversational knowledge management and group collaboration. *Communication of the Association for Information Systems*, 13 (2004), 265-289.
- 88 Wagner, C. and Majchrzak, A. Enabling customer-centricity using wikis and the wiki way. *Journal of Management Information Systems*, 23, 3 (2007), 17-43.

- 89 Wall, V. and Nolan, L. Perceptions of inequity, satisfaction, and conflict in task-oriented groups. *Human Relations*, 39 (1986), 1033-1052.
- 90 Wang, R. Y. and Strong, D. M. Beyond accuracy: What data quality means to data consumers. *Journal of Management Information Systems*, 12, 4 (1996), 5-33.
- 91 Weber, S. *The success of open source*. Cambridge: Harvard University Press, 2004.
- 92 Wellman, B. S. An electronic group is virtually a social network. In S. Kiesler (ed.), *Culture of the Internet*. Mahwah: Psychology Press, 1997, pp. 179-206.
- 93 Williams, R. L. and Cothrel, J. Four smart ways to run online communities. *Sloan Management Review*, 41, 4 (2000), 81-91.
- 94 Williams, K. and O'Reilly, C. The complexity of diversity: A review of forty years of research. In B. Staw and R. Sutton (eds.), *Research in organizational behavior*. Greenwich: JAI Press, 1998, pp. 77-140.
- 95 Wolfe, C. J. and Murthy, U. S. Negotiation support systems in budget negotiations: An experimental analysis. *Journal of Management Information Systems*, 22, 3 (2005), 351-381.
- 96 Wong, S. Task knowledge overlap and knowledge variety: The role of advice network structures and impact on group effectiveness. *Journal of Organizational Behavior*, 29 (2008), 591-614.
- 97 Woodman, R. W.; Sawyer, J. E.; and Griffin, R. W. Toward a theory of organizational creativity. *The Academy of Management Review*, 18, 2 (1993), 293-321.

Appendix A – Procedure for Analyzing Task Conflict in Wikipedia Discussion Pages

- Each article’s discussion page was analyzed, independently, by three research assistants, under the supervision of the authors. Assessors processed the documents in a random order.
- Since the task at hand is the authoring of an encyclopedic entry, task-related conflicts are disagreements between contributors regarding the contents of the article being authored. The product of the authoring process should be an encyclopedic entry that is accurate, comprehensive, unbiased, and clear, and thus conflicts often relate to the accuracy, coverage, objectivity, and clarity of the article. Some examples of task-conflict related statements found in the discussion pages include: “Jim was born in Argentina in 1954, and not in Belgium” or “This information should be removed, as it is not directly related to the topic of the article”

Once an assessor completed reading the article’s discussion page, she rated his perception of conflict, with respect to the following statements:

- “Participants often have conflicting opinions about what should be included in the article.”
- “Participants frequently have disagreements about what information the article should include.”
- “There exists substantial conflict of ideas among participants.”

Figures

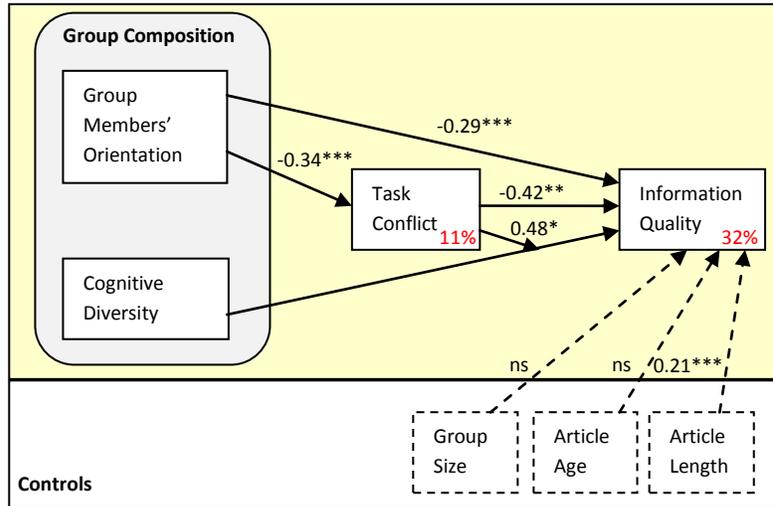


Figure 1: Results of PLS analysis; values on arrows represent path significance; ‘*’ indicates significance of $p < 0.05$, ‘**’ signifies $p < 0.01$, ‘***’ indicates significance of $p < 0.001$, and “ns” indicates a non-significant path. R² percentages in red.

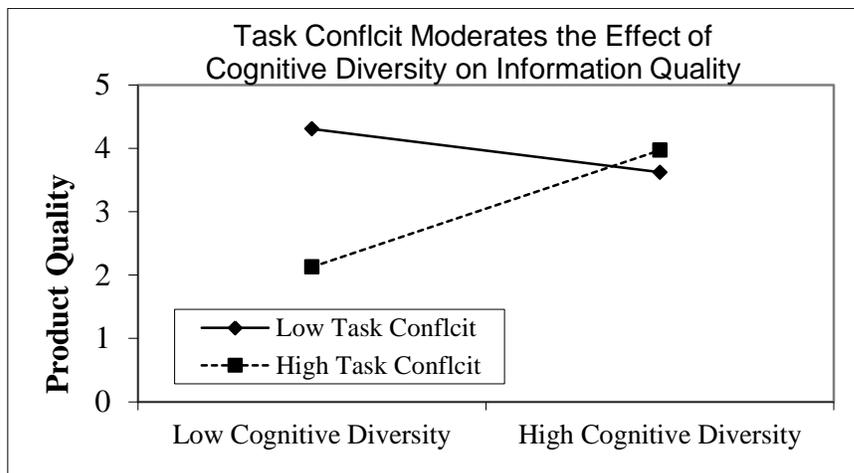


Figure 2: Task conflict moderates the effect of cognitive diversity on information quality.

Tables

Table 1: the characteristics of content- and administrative-oriented members

Member's Orientation	Commitment and identification with the community	Engagement and participation in community	Pattern of Contribution	Roles and responsibilities
Content	Low	Low	Centered on few items	Basic
Administrative	High	High	Distributed across many items	Volunteer for additional responsibilities

Table 2: descriptive statistics

Construct	Measure	Mean	STD
Group Members' Orientation	Percent of administrators	0.15	0.113
	Average of members' Wikipedia edits	6217.1	5204.2
	Ratio of members' Wikipedia edits by article edits	3552.9	3084.8
	Average of members' entropy in Wikipedia article edits	4.89	1.32
Cognitive Diversity	The sparsity of the editors to external articles matrix	0.92	0.077
Task Conflict	"Participants often have conflicting opinions about what should be included in the article"	2.1	1.24
	"Participants frequently have disagreements about what information the article should include"	2.1	1.23
	"There exists substantial conflict of ideas among participants"	2.0	1.12
Group Size	Number of unique editors who have contributed to the article	49.2	70.4
Article Length	Number of words in the article	901.6	825.2
Article Age	Days since the inception of the article	812.2	462.3
Article Activity	Number of edits made to the article since its inception	90.9	125.0
Information Quality	Raters' consensus on the quality of article	4.4	1.16

Table 3: Square-root of AVE (in bold, on diagonal) and correlation between the latent constructs.

Constructs	<i>CD</i>	<i>GMO</i>	<i>TC</i>	<i>TC x CD</i>	<i>GS</i>	<i>AL</i>	<i>AAg</i>	<i>AAc</i>	<i>IQ</i>
<i>Cognitive Diversity (CD)</i>	1								
<i>Group Members' Orientation (GMO)</i>	-0.41	0.89							
<i>Task Conflict (TC)</i>	0.49	-0.34	0.99						
<i>TC x CD</i>	-0.48	0.14	0.38	0.99					
<i>Group Size (GS)</i>	0.49	-0.31	0.63	0.32	1				
<i>Article Length (AL)</i>	0.20	-0.15	0.44	0.27	0.48	1			
<i>Article Age (AAg)</i>	0.37	0.15	0.30	0.08	0.47	0.18	1		
<i>Article Activity (AAc)</i>	0.49	-0.33	0.67	0.35	0.98	0.50	0.44	1	
<i>Information Quality (IQ)</i>	0.14	-0.30	0.24	0.26	0.45	0.36	0.15	0.46	1