

The Hidden Order of Wikipedia

Fernanda B. Viégas, Martin Wattenberg, and Matthew M. McKeon

Visual Communication Lab, IBM Research, 1 Rogers St,
Cambridge, MA 02142, USA
{viégasf, mwatten, mmmckeon}@us.ibm.com

Abstract. We examine the procedural side of Wikipedia, the well-known internet encyclopedia. Despite the lack of structure in the underlying wiki technology, users abide by hundreds of rules and follow well-defined processes. Our case study is the Featured Article (FA) process, one of the best established procedures on the site. We analyze the FA process through the theoretical framework of commons governance, and demonstrate how this process blends elements of traditional workflow with peer production. We conclude that rather than encouraging anarchy, many aspects of wiki technology lend themselves to the collective creation of formalized process and policy.

Keywords: Wikipedia, Governance, Commons, Peer Production

1 Introduction

Wikipedia, the online encyclopedia that anyone can edit, has become one of the most visited sites on the web. A common question is how such apparently useful content can be generated by an army of distributed volunteer editors. This paper discusses part of the answer: despite the seeming potential for anarchy or chaos, a sophisticated set of processes have emerged.

Every day the Wikipedia front page presents a “Featured Article” (FA). Consider the trajectory of one such article, on AIDS, which was featured on June 15 2006. Before appearing on the front page, the article underwent a lengthy process of peer review. It was nominated as a Featured Article Candidate (FAC) on 20 March 2006. Before nomination, it had gone through a separate peer review to help improve its quality. The article was also part of the “Medicine Collaboration of the Week” project, where members focus their attention on a given medicine-related article per week. The FAC review process itself involved 18 different users and amassed 61 posts. In these posts there were references to seven of Wikipedia’s guidelines and the entire review process lasted 20 days.

Such a complex and bureaucratic process runs counter to naïve depictions of Wikipedia as an anarchic space. The site boasts myriad guidelines, policies and rules. Moreover, a series of formal processes, of which the FA procedure is a prime example, are starting to materialize. Analyzing the organizational principles behind these emerging processes can help us better understand the inner workings of Wikipedia. The emergence of these processes is, we believe, just as interesting—and “magical”—as the emergence of high-quality articles.

The paper starts with a review of related academic work on online communities. This review leads us to the work of two scholars who have framed our analysis of process in Wikipedia: Yochai Benkler's study of online peer production [1,2] and Elinor Ostrom's review of collective action and governance in offline communities [14, 15]. After laying out this intellectual framework, we plunge into a detailed investigation of the FA process in Wikipedia: how it works, how it relates to Wikipedia's policies and guidelines, and what it tells us about formal processes on the site. We finish with a discussion of how the structure of the FA process relates to Benkler's and Ostrom's principles. We conclude that rather than encouraging anarchy, many aspects of wiki technology lend themselves to the group creation of workflows and process.

2 Related Work

In the 1980s and earlier, online communities consisted mainly of conversation. People came together in Usenet, chatrooms, IRC channels, and even MUDs primarily to talk to each other. With the exception of MUDs, interaction was exclusively conversational and scholarship on these environments reflected the focus on conversation [4, 7, 13].

Online communities today include social networking sites, wikis, and social bookmarking tools. While conversation remains important, the production of a variety of goods has become a vital aspect of these communities. Code, encyclopedic entries, massive websites, and even game economies exist today. In response, online communities scholarship has expanded to encompass new inquiry areas such as economics and law [10].

To examine the governance structure of Wikipedia, this paper draws on the literature of regulation in commons-based communities. Our main sources are Yochai Benkler's work on commons-based peer-production and Elinor Ostrom's work on commons-based governance. After an overview of this theory, we then focus on the FA process in Wikipedia and its relation to Ostrom and Benkler's work.

2.1 Commons-Based Peer Production

One of the best-studied examples of online collaborative production is open source software. Researchers have examined both the development process [12] and the incentive structures [9]. Of particular interest is a framework proposed by Yochai Benkler encompassing open-source development, Wikipedia, and several other online systems. Benkler suggests these systems represent a new form of economic organization, distinct from either firms or markets: commons-based peer-production [1]. Unlike other organizational methods—such as the market and the firm—peer production depends on individual action that is self-selected and decentralized rather than hierarchically assigned. Individuals make their own choices with regard to resources managed as a commons.

Benkler defines two evolutionary phases that successful commons-based communities typically go through:

1. *Creating content (utterance):*

This is the initial phase where large, complex tasks are broken into small, independent modules. This phase is marked by providing contributors with a wealth of tasks that can be achieved individually, in uncoordinated fashion.

2. *Quality control (relevance/accreditation):*

This second phase is characterized by a concerted effort on quality assurance. How can we know that the content produced by widely dispersed individuals is not nonsense? In this phase, the community must define standards and create low-cost quality control mechanisms.

Scholars are starting to investigate the nature of work coordination and quality assurance in Wikipedia. Stvilia et al. investigated how Wikipedians improve the quality of entries through discussions in Talk pages [18]. Viégas and colleagues have examined the role of Talk pages in group coordination and policy enforcement [20].

2.2 **Governing the Commons Offline**

Online communities are not the only place where one finds commons-based communities. In fact, the challenges of commons-based governance are not new. There is a broad literature about the evolution of institutions for collective action in the offline world, where communities have had to self-organize and govern for millennia. Elinor Ostrom's work [14,15] analyzes the principles behind successful, self-governed common-pool resources communities (CPRs). She has looked at communal tenure in a variety of settings, including high mountain forests in Japan, commons-based irrigation institutions in Spain and the Philippines, and inshore fisheries in Turkey and Sri Lanka. For centuries, these communities of farmers, villagers, and fishermen have successfully found ways to manage shared natural resources—forests, rivers, fisheries, timber—without relying on centralized authority.

Some challenges faced by these offline, self-governed communities are similar to the challenges faced today by their online counterparts: creation of rules, monitoring mechanisms, arbitration, and conflict resolution. Ostrom proposes a list of eight organizational principles found in long-enduring CPR institutions. As a first step, we focus on four principles which seem natural to map to the online space. A detailed consideration of the other principles is an important area for future research.

1. Congruence between rules and local conditions: instead of relying on “one-size-fits-all” regulation, rules must be intimately associated with the particularities of the resources they regulate. For instance, a community of farmers who depend on a river for irrigating their crops will need to devise rules that fit the particular geographical profile of their region and river as opposed to relying on some “generic” set of rules for irrigation.

2. Collective-choice arrangements: most individuals affected by the operational rules should be able to participate in modifying these rules and the cost of altering rules should be kept low.

3. Monitoring: Individuals who monitor the commons should be accountable to the rest of the community.

4. Conflict-resolution mechanisms: community members should have rapid access to low-cost local arenas to resolve conflicts.

Whereas Benkler's work addresses the importance of coordination and information flow in online peer-production, Ostrom's principles show us how commons-based communities can be successful in self-organization and self-governance. Her principles give us a framework with which to examine and situate coordination processes in Wikipedia.

3 Case Study: Featured Articles

To examine the interplay of rules, processes, and governance in Wikipedia in depth, we concentrated on one of the best established and most visible processes on the English site: the FA process [22]. As mentioned above, an FA, is an article that appears prominently on the main Wikipedia page. There is only one FA at a time, changing once a day. Selecting these articles is a delicate matter: quality is important, since they represent the public face of Wikipedia. At the same time, many more articles are suggested for FAs than can be accommodated. Choosing which articles should be featured is a challenge in collective action, and it turns out that a process has evolved in response.

We gathered information on the FA process through several avenues. An important aspect of Wikipedia is that procedures and the guidelines that drive them are described in detail on publicly accessible pages. Furthermore, much of the discussion surrounding the creation of the guidelines is available via so-called “Talk” pages. Thus the point of departure for our investigation was a careful reading of FAs, FA reviews, guideline pages, and discussion pages attached to guidelines.

To augment this examination, we conducted an extended interview, via email and telephone, with one of the key players in the FA process, Mark Pellegrini, the director of the front-page Featured Articles. This interview was useful both for confirming facts we had learned from reading the history of the process as well as providing some of the organizational “backstory” of the process.

Over the years, the standards for promoting an article to FA status have increased dramatically. In the beginning, for instance, there was no requirement that an article must contain inline citations. Moreover, requirements for topic comprehensiveness have become stricter. In fact, the criteria have evolved so much that over 200 of the early FAs have been demoted because they do not meet current FA criteria.

Before an article can be promoted to FA status, it needs to be nominated as a Featured Article Candidate (FAC). Anyone can nominate an article as an FAC. Nominations are public and nominators are expected to make an effort to address any objections that editors raise during the review process (see “FAC review” in figure 1). For instance, if an editor objects to the prose style of the lead section of the article, the nominator is expected to rewrite it. When nominators have worked on the article prior to nomination, they are supposed to mark it a “self-nomination.”

A nomination summarizes the state of the article, For example:

Daniel Boone: Self-nomination. Listed as a "good article", assessed as "A-class" by Wikipedia WikiProject Military History, has gone through a couple of peer reviews. The article is based on the major 20th century biographies, with points of disagreement between historians noted in the text or footnotes, especially regarding the issue of history versus folklore, a central concern in Boone historiography. All comments are welcome; hope you enjoy reading it. — Joan 16:43, 18 September 2006 (UTC)

Supporting and objecting

For a nomination to be promoted to FA status, consensus must be reached that the article meets the FA criteria. Anyone is allowed to participate in the process of reviewing an FAC and votes of support or opposition need to be backed by explicit

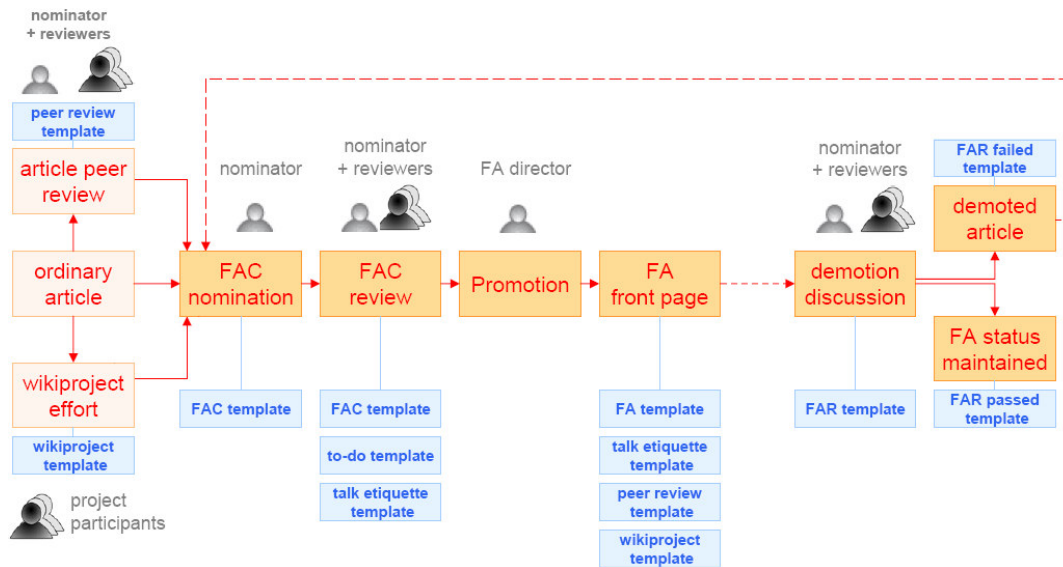


Fig. 1: Diagram of steps in the Featured Article process. Before an article is nominated as an FAC, it tends to undergo significant editing—it is not uncommon for articles to have gone through a separate peer review or have been the focus of a Wikiproject. Once the article is nominated, it enters the FA pipeline, which includes a review and, hopefully, promotion. All along, templates communicate the status of the article to contributors and readers.

reasoning. Objections have to be actionable, in the sense that they have a clear way of being addressed; if nothing can be done in principle to address the objection, the FA Director may ignore it. An example of an actionable objection is:

Nagorno-Karabakh War → *Object: massively undercited; many large portions of text—and even direct quotes!—have no citations. More generally, I'm concerned that almost all the references seem to be newspaper articles, even though a number of books dealing with the topic (including those listed as "Further reading", for example) are not used as sources.* -Jason 19:07, 16 September 2006 (UTC)

The FA Director determines whether there is consensus on promoting or rejecting an FAC. If actionable objections have not been resolved or consensus for promotion has not been reached, a nomination may be removed from the list and archived. The FA Director determines the timing of the process for each nomination. The FAC process generally takes at least 5 days.

Articles can lose their FA status over time (“demotion discussion,” Figure 1). While this may seem natural since entries change over time, articles usually lose FA standing not because of edits but because they the FA criteria have become stricter. FAs are demoted through a consensus derived through discussion on the FA Removal Candidates page. A user has to nominate the entry for demotion, which initiates a review process.

That FA entries can be downgraded because of sub-standard quality is a testament to the evolution of the FA process. As of September 2006, 238 articles had lost featured status, while eight had been re-promoted.

Several automation tools have been created for the FA process. One looks for common problems of syntax and style, for example the use of “weasel words” and automatically creates a to-do list for the page. A second, the cite.php module, created in

2005, helps with various aspects of creating inline citations. Before the module existed, creating such citations was awkward; the introduction of the module has been mentioned as an enabling factor for the current strict requirements for citations.

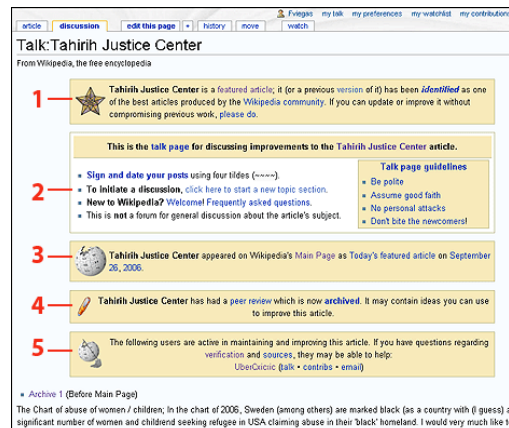


Fig. 2: Screenshot of templates at the top of the discussion page of an FA.

Workflow Markers: Templates

The FAC process can be thought of as the workflow necessary to promote an article to FA status. There is a series of steps that the article must undergo and a sequence of tasks involved in the procedure (Fig. 1). Keeping track of this progression can be hard. To solve this problem, Wikipedians have devised a way to communicate the status of an article: *templates*. A template is a piece of wiki code that creates a visual marker—often a text box with a different background color from that of normal text (Fig. 2)—usually placed near the top of a page. In the context of the FA process, templates often alert Wikipedia readers and contributors to the current status of an article. For instance, templates can include navigation aids, warnings that a page is currently being featured on the home page, or that content is sub-standard.

A page can have several templates and any contributor is free to add templates to pages. In fact, templates are so important that they have their own namespace and are used on pages throughout the encyclopedia. Templates that provide information only of service to editors belong on an article's talk page. These are the kinds of templates that aid the FAC review process.

Editors use templates to keep contributors informed of the current status of the entry. In figure 2, for example, template (1) says this article is an FA, (2) lists Talk page guidelines and briefly explains Talk page editing etiquette, (3) documents the date when the article appeared on Wikipedia's Main Page, (4) indicates that the article has had a peer review which is now archived, and (5) lists users who are active in maintaining and improving this article.

A special kind of template is particularly useful for editors during the FAC process: the to-do list. As its name implies, this template lists improvements that are

suggested for the article. The list is maintained by editors, writers, reviewers or readers as a way to focus collaborative efforts

Users can add templates to their personal watchlists, so they are notified any time that template is added to a new page. Articles with a to-do template, for instance, are automatically inserted into the list of articles with To do's, which attracts additional editors. When a to-do is finished, editors strike it out to mark the progress.

4 Discussion

The FA process consists of a well-documented series of steps organized around an artifact—an article—which is modified, approved, or rejected. The process for moving through individual steps is rule-bound, guided by a large set of written policies. In fact, the FA endeavor starts to sound very much like a modern-day, enterprise workflow process. It is not, however. Here we discuss the FA process in light of Benkler's and Ostrom's work.

Benkler: Peer-production and Peer-Process

Two aspects of the FA process are unusual. First, several roles in the process are filled by crowds of self-identified individuals. The editors of the article, the reviewers, and the people who vote on whether the article meets FA criteria are volunteers and there is no preset limit on the number who may participate. The second unusual aspect is the non-hierarchical flow of information, where some people signal that work is needed—through the use of templates—and other people pick up the signal and act on it. While there is a FA Director, he relies completely on volunteers and is not a “boss.”

This arrangement resonates with Benkler's claims about the first phase of commons-based peer production and the importance of breaking large, complex tasks are into small, independent modules. In this case, the independent modules are the small edits and votes needed to move an article along in the FA promotion process. As soon as someone adds an “FAC template” to a page, volunteers will find that page and may decide to review the article. These individuals may have no previous connection with the article, and they may never look at it again in the future, but they will spend time reviewing it.

At the same time, however, the FA process does not belong entirely in Benkler's first phase because its *raison d'être* is quality assurance. So here is an interesting hybrid of both of Benkler's phases: a process that coordinates individuals' efforts around quality assurance (phase II), while doing so in a distributed manner that relies on independent modules—“five-minute increments of human attention” (phase I). Part of the FA success is likely due to the fact that individuals can easily step in and out of the process at any point. It is a “peer-process:” a completely distributed, yet coordinated and formalized, procedure.

Ostrom: Policy and Self-Governance

Another key reason why the FA process runs smoothly is Wikipedia's extensive body of rules and guidelines (as of September 2006, 75 general guidelines and 119 style guidelines covering text presentation and formatting). The policies were written

by the community to address a set of problems that is common to all efforts to organize collective action: creation of institutions, monitoring mechanisms, arbitration, and conflict resolution. These are exactly the challenges faced by the self-governing communities studied by Elinor Ostrom that succeeded in managing natural resources . There is an impressive degree of overlap between what happens on Wikipedia and the design principles that Ostrom extracted from. offline, communities.

That Wikipedians have independently arrived at some of the same governance answers as in offline communities suggests some of these principles are universal. Conversely, an analysis of which principles do not hold in Wikipedia may inform us about what is particular to online self-governance. Here we describe how four of Ostrom's principles translate directly to the context of Wikipedia: congruence between rules and local conditions, collective-choice arrangements, low-cost monitoring, and conflict-resolution mechanisms. We then add a new hypothesis, that the persistent, public nature of work and debate on Wikipedia is a key to the success of peer process and briefly discuss the role of templates as workflow markers.

Congruence between rules and local conditions: There is close interaction between rule shaping and what is happening “in the field” in Wikipedia. An example of this interaction is the adaptation of the article size rule to the new reality of the FA requirements. Originally, FA criteria did not require references and citations explicitly. Adding properly linked citations was an awkward manual process. Now, the standard FA must have inline citations and a comprehensive coverage of its subject matter. This change has been accompanied both by accommodations in code (the cite.php module that eases the manual labor) and in rules: page size now refers only to the prose in the main body of the article, disregarding references and citations.

Collective-choice arrangements: This principle means that most individuals affected by operational rules can participate in modifying the rules. This is true of Wikipedia where rules are publicly discussed and established. Anyone can participate in the debate about rules and policies. Moreover, the costs of changing rules is low, at least from a technical perspective—anyone can post suggestions to a talk page and make the case for why a given rule needs to change. The low barriers for participation mean that regulations are not set in stone and can be adapted to better fit the intention of the community.

An important convention for collective choice in Wikipedia is polling. Indeed, the culture of Wikipedia seems to encourage polling as a means of building consensus and voting as a way of making a choice. This mechanism has been important in the development of the FA review procedure, starting with the original vote to delegate authority to the FA Director to make changes to the home page.

Monitoring: Ostrom posits that a well-governed community needs to have low-cost monitoring capabilities to prevent free riding or other antisocial, negative behavior. While this concern certainly holds for Wikipedia in general and the FA process in particular, the element of online asynchronous work done by ad hoc volunteers adds additional considerations to the notion of monitoring.

Consider the role that users' watchlists play in monitoring activity. Not only do they allow editors to quickly address harm done to the site, they also provides users

with a way of organizing editing activity. The interaction between watchlists and templates transforms a monitoring mechanism into a tool for work coordination.

For instance, in the FA process, the templates allow participants to see at a glance the status of an article, thus providing visibility into the progress of the process. An additional technical point is that by adding templates to a watchlist a user is able to see all pages newly tagged with that template. This fact is crucial to the functioning of the FA process: merely by tagging an article with the appropriate template, a user can attract the attention of a crowd of volunteers willing to participate in the process.

Easily-accessible, persistent public documentation and conflict management: the persistent and public nature of transactions on the site helps users coordinate actions and resolve disputes. In fact, having accessible records is so crucial for the community that most communication between users happens publicly on the site [16]. Although Wikipedians use mailing lists, they are encouraged to keep most communication within the confines of the public site. This approach creates a transparent system of record keeping that is easy to refer and link to. In turn, these easy-to-access records are invaluable for the evolution of governance. The FA process provides several examples of persistent archives. During editing activity, our findings show that it is common for editors to refer to a many different guideline pages. These guidelines provide common ground for participants, help resolve arguments, and ensure consistency across instances of the process.

A second simple example is given by the FA Director [16]. One commonly requested change to the process is to protect a page on the day it is a FA, to defend against the expected onslaught of vandalism. After much debate, it was decided not to take this step, yet new users consistently ask for it to be implemented. The FA Director reports that he simply posted the reasoning behind the debate on his user page, creating a document he could refer people to when the question arises.

Technology: Two pieces of technology conspire to make the sort of peer-process we find in Wikipedia possible. On one hand, the persistence of wiki records means that all steps of the process are continuously documents and available. On the other hand, Wikipedians' creative use of templates means that tight coordination of activity does not impose any type of workflow overhead on participants.

5 Conclusion

The vast number of policies in Wikipedia and the existence of robust, formal processes such as FA, indicate that governance is a thriving aspect of this community. The fact that these policies and processes have been devised and modified over time according to a set of collective-choice rules makes Wikipedia a fascinating example of self-governing institutions. In fact, Wikipedia's formalized processes, such as FA, seem to share several of the design principles found by Ostrom in offline, self-governed communities around the world [15].

First, Wikipedia's rules are tightly bound to particular technical aspects of the site and are therefore "localized," instead of generic solutions. Second, they are collective-choice arrangements where everyone may participate in modifying the

rules. Third, monitoring the actions of others is facilitated by the technology available. In fact, we identify Wikipedians' use of templates as being one of the driving technological factors of the success of the FA process. Finally, in contrast with many offline communities, Wikipedia's records are persistent, public, and easily available online. We believe this is the second technological element driving both the creation and the adoption of norms and guidelines on the site.

A large part of the increase in coordination and regulation efforts in Wikipedia is due to the need of defining quality standards and assuring quality control in entries. The FA process is the poster child of this endeavor. We find that this process represents a hybrid of Benkler's two evolutionary phases in successful commons-based communities: the FA process directly addresses quality assurance at the same time that it is structured to allow complex coordination tasks to be broken into small, independent modules. In other words, FAs ensure quality using peer-production mechanisms, and this characteristic is likely a key to their success in Wikipedia.

References

1. Benkler, Y. Coase's Penguin, or, Linux and The Nature of the Firm. *The Yale Law Journal*, Vol 12, N 3. December 2002.
2. Benkler, Y. *The Wealth of Networks: How Social Production Transforms Markets and Freedom*. Yale Press, 2006.
3. Bryant, S., Forte, A., Bruckman, A. Becoming Wikipedian: Transformation of Participation in a Collaborative Online Encyclopedia. In *Proceedings of GROUP 2005*.
4. Cherny, L. *Conversation and Community*. CSLI Pub., 1999.
5. Emigh, W., and Herring, S. Collaborative authoring on the Web: A genre analysis of online encyclopedias. In *Proceedings of HICSS-38*, 2005.
6. Forte, A., Bruckman, A. Why do People Write for Wikipedia? Incentives to Contribute to Open-Content Publishing. *GROUP 05*.
7. Herring, S. Interactional coherence in CMC. *HICSS 32*, 1999.
8. Holloway, T., Bozicevic, M., Börner, K. Analyzing and Visualizing the Semantic Coverage of Wikipedia and Its Authors. Submitted to *Complexity*, 2005.
9. Lakhani, K., & Wolf, R. Why Hackers Do What They Do: Understanding Motivation and Effort in Free/Open Source Software Projects. In *Perspectives on Free and Open Source Software*, MIT Press: 2005.
10. Lessig, L. The Law of the Horse: What Cyberlaw Might Teach, 113 *Harvard Law Review* 501, 1999.
11. Lih, A. Wikipedia as Participatory Journalism: Reliable Sources? Metrics for Evaluating Collaborative Media as a News Source. *Proceedings of the Fifth International Symposium on Online Journalism*, 2004.
12. Mockus, A., Fielding, R., & Herbsleb, J. Two Case Studies of Open Source Software Development: Apache and Mozilla. *ACM Trans. Software Engineering and Methodology*, 11(3), 309-346, 2002.
13. Nonnecke, B. & Preece, J. Lurker Demographics: Counting the Silent. *CHI 2000*.
14. Ostrom, E. *Governing the Commons: The Evolution of Institutions for Collective Action*. New York: Cambridge University Press, 1990.
15. Ostrom, E. Collective Action and the Evolution of Social Norms. *Journal of Economic Perspectives*, Vol. 14, No. 3, pp. 137-158, 2000.
16. Pellegrini, M. (Featured Article Director), personal communication, September 2006.
17. Sartwell, Crispin. Wikipedia: See 'Information,' 'Amazing, Anarchy.' *Los Angeles Times*, May 4, 2005
18. Stvilia, B., Twidale, M., Gasser, L., Smith, L. Information Quality Discussions in Wikipedia. Technical Report ISRN UIUCLIS--2005/2+CSCW, 2005.
19. Viégas, F., Wattenberg, M., & Dave, K. Studying Cooperation and Conflict between Authors with history flow Visualizations. In *Proceedings of SIGCHI 2004*.
20. Viégas, F., Wattenberg, M., Kriss, J., & van Ham, F. Talk Before You Type: Coordination in Wikipedia. *HICSS-40*.
21. Wales, J. Wikimania Keynote Address. Cambridge, MA 2006.
22. http://en.wikipedia.org/wiki/Wikipedia:Featured_articles