

Scaling Consensus: Increasing Decentralization in Wikipedia Governance

Andrea Forte and Amy Bruckman

GVU Center, College of Computing, Georgia Institute of Technology
{aforte, asb}@cc.gatech.edu

Abstract

How does “self-governance” happen in Wikipedia? Through in-depth interviews with eleven individuals who have held a variety of responsibilities in the English Wikipedia, we obtained rich descriptions of how various forces produce and regulate social structures on the site. Our analysis describes Wikipedia as an organization with highly refined policies, norms, and a technological architecture that supports organizational ideals of consensus building and discussion. We describe how governance in the site is becoming increasingly decentralized as the community grows and how this is predicted by theories of commons-based governance developed in offline contexts. The trend of decentralization is noticeable with respect to both content-related decision making processes and social structures that regulate user behavior.

1. The Mechanisms of Self-Organization

Should a picture of a big, hairy tarantula appear in an encyclopedia article about arachnophobia? Does it illustrate the point, or just frighten potential readers? Reasonable people might disagree on this question. In a freely editable site like Wikipedia, anyone can add the photo, and someone else can remove it. And someone can add it back, and the process continues. In fact that's exactly what happened on Wikipedia from July 2006 through January 2007, until the community finally settled on including a cartoon about fear of spiders instead of a photo of the real thing. But not before hours had been spent on the topic by several editors. In fact, at the time this paper was written, more effort (roughly 6000 words of discussion) has been put into the debate about the photo than into discussion about the rest of the article (roughly 1500 words). If something as simple as a spider photo can generate this much controversy, how does Wikipedia deal with issues like the content of articles on Palestine or global warming? It seems almost miraculous that editors don't get permanently bogged down in “edit wars” and that real work gets done at all. Yet somehow, despite common episodes of controversy and disagreement among authors, it does.

In 2007, the Virginia Tech shootings led to a debate about whether a Wikipedia page should exist for each victim. How can such an issue be resolved? As we will see, policies and social norms have evolved over many years to guide the production of biographical content on the site and these were used to make decisions regarding Virginia Tech victims. But how are such policies created and how have they changed over time? How are policies interpreted and enforced? This paper investigates how the English-language Wikipedia is currently governed and how this process has evolved over time using biography-related policies as a central example.

Wikipedia governance is always changing, adapting to new challenges. One of these challenges has been the substantial growth and flood of new editors that has accompanied the site's popularity. We describe the ways in which governance in the site is becoming increasingly decentralized over time as the community responds to the challenge of growth and why the form that decentralization takes is consistent with literature on commons-based resource management. The trend of decentralization is noticeable with respect to both content-related decision making processes and structures that regulate user behavior. We build on the work of Viegas et. al. in using Ostrom's principles of self-organized communities to understand governance mechanisms on the site [9, 10, 18].

2. Related Work

For over a decade, online communities scholars have revealed intricate systems of social stratification and hierarchy, identity formation, virtual economies and approaches to dispute resolution that help regulate behavior in online social spaces [6, 11, 14]. Recently, peer-production models have become important features both of the global economic landscape and of life on the Internet. To understand how people organize themselves around the production of goods in volunteer organizations, scholars have examined social organization in online communities that generate content [1, 2, 5, 15-17].

When we talk about governance, what do we mean? Governance generally refers to a system for organizing the rules that regulate people's behavior in a particular

place; however, formal systems of governance and law are not the only things that regulate behavior. In *Code*, Larry Lessig offers an explanation of what things regulate individuals' actions: 1) formal governance/law 2) social norms 3) the market and 4) architecture/technology—in the case of the Internet, code [7]. It is important to understand that these regulators interact and all four are constantly at work constraining the behavior of Wikipedians even as they create, interpret, follow and enforce policies. External regulators, such as the laws and cultures of nations in which Wikipedians live, also constrain behavior.

In this paper, we build on Viegas et. al.'s "Hidden Order of Wikipedia" [18]. Viegas et. al. borrowed from the literature on self-organizing communities to analyze the procedures associated with selecting Featured Articles (FA) in Wikipedia (2007). Their analysis was centered on four of Elinor Ostrom's eight design principles for self-organizing communities that manage natural resources [9]:

1. Congruence between rules and local conditions
2. Collective-choice arrangements
3. Monitoring
4. Conflict-resolution mechanisms

These four principles are highly salient for understanding Wikipedia governance both in the particular case of FA selection and more generally. Our analysis explores a broad swatch of governance-related activity in the site. As we describe the trend of increasing decentralization in Wikipedia, we will also point out where Ostrom's remaining four principles are embodied in Wikipedia governance:

5. Clearly defined community relationships
6. Graduated sanctions
7. Local enforcement of local rules
8. Multiple layers of nested enterprises

Although Wikipedia is not a natural resource, it is a valuable one. Community members strive to manage Wikipedia namespace as its value increases with the site's popularity and the consequences of misuse become more dire. It should be noted that, like many natural resources, namespace is consumable; only one article exists on any given topic in Wikipedia and improper use of namespace, though recoverable, is a loss. Of course, differences do exist between online and physical resources. For example, in communities studied by Ostrom, defining membership boundaries is important because members reap the rewards from resources being managed. In the case of Wikipedia, this distinction does not have the same meaning since millions of people outside the community benefit directly and immediately from the efforts of its membership.

Still, we believe that Ostrom's work is a natural fit for understanding processes of governance on Wikipedia because the community is not only managing a resource,

it is striving to encourage collaboration and cooperation among volunteers. Ostrom's work is built on the proposition that the evolution of social norms within a community is a more effective means of accomplishing cooperation than the imposition of external rules. She notes not only that social norms are equally effective at "generating cooperative behavior" but also that "moreover, norms seem to have a certain staying power in encouraging growth of the desire for cooperative behavior over time, while cooperation enforced by externally imposed rules can disappear very quickly" p.147 [9]. In Wikipedia, governance is less akin to rule imposition by external authority than to constant reform and refinement of social norms within the community.

In his chapter on "governance" of synthetic worlds, Castronova notes that one clear obstruction to community-organized governance online is when users are not equipped with the technical powers required to perform acts of governance [4]. The idea is not new. Morningstar and Farmer argued as early as 1990 for developers to relinquish control, and that users "should be able to materially effect each other in ways that went beyond simply talking, ways that required real moral choices to be made by the participants," and recommended that "a virtual world need not be set up with a 'default' government, but can instead evolve one as needed" [8]. The potential for sophisticated, community-generated social norms and governance mechanisms is partly a designed feature of the technological architecture in which the community grows. In other words, artifacts have politics (see [20]) and code influences the development of policy. In Wikipedia, we find that code, policy, and social norms all support the wide distribution of power to govern.

2. Method: Interpreting Wiki Activity

Wikipedia did not arise spontaneously, it arose through people interacting and, as a result of that interaction, finding ways that worked. - Interviewee 3 (13)

In order to understand what regulates behavior in Wikipedia, we interviewed individuals who had experienced those regulating influences first hand. Our approach is influenced by phenomenological approaches to sociology [12, 19]. We believe that understanding a social space involves gaining access to the experiences and interpretations of people who live in that world. Phenomenological methods for investigating social spaces include interviewing and participant observation. (See [13] for a discussion of interviewing as a phenomenological method.) Although our formal analysis here is concerned with data collected through interviews, we have spent hundreds of hours over several years reading Wikipedia community mailing lists and talk pages discussions, following articles, getting to

know members of the Wikipedia community, and building our own extensions to MediaWiki—the software on which Wikipedia is built. These activities have honed our ability to act more reliably as instruments of interpretation and allow us to make sense of participants’ stories.

Interview-based research is about the participants; what one can learn from them is constrained by their experience. Sampling methods, then, are a critical feature of this study. We used a layered approach to strategically recruit participants who could provide us with insights about specific issues that we wished to understand better. We began by soliciting five interviews from long-standing, central members of the community who told us their stories about how norms, policies, social roles, and software have changed over time. When themes began arising over and over in these interviews, we followed up by recruiting participants who had been involved in particular issues and policies. In all, we interviewed ten individuals by telephone and one via email, all of whom had been involved in the site in many different capacities (See Table 2 in the next section for a partial list). These interviews resulted in roughly nine and a half hours of audio and 100 pages of transcription.

In traditional social science research, participants are anonymized to protect their privacy. However, this becomes problematic in cases where participants have made significant contributions for which they deserve credit [3]. Additionally, public figures like Jimmy Wales are impossible to anonymize. With Wales’ consent and permission from our Institutional Review Board (IRB), we identify him by name and use “light disguise” (Ibid.) for other prominent Wikipedia figures interviewed.

3. A Lexicon of Wikipedian Social Life

Before launching into a thick description of regulatory mechanisms in Wikipedia, it is helpful to understand some of the different kinds of social roles that make up the community. This initial, impoverished description of Wikipedian roles will be elaborated in the following sections. Power relationships that exist on the site are complex and subtle; however, we have identified three basic user states that affect in fundamental ways the kind of social authority individuals can exert in the English language site: unregistered user, registered user, Arbitration Committee member.

Unregistered users can exert little individual influence in shaping policy and establishing norms, but en masse they represent an important part of the context in which day-to-day operations take place. In most cases, unregistered users have the ability to edit the encyclopedia freely but one interviewee noted that their ability to influence content is weaker than registered

users, “You can’t track anything. The community views their edits to be particularly suspicious” (I7).

Registered Users includes everybody else on the site; we will use “regular users” to refer to registered users who hold no special technical powers. It is impossible to describe all the possible roles that regular users can play and the nuanced forms of authority they may hold. Power relationships are difficult to define. “It’s such a strange place because the structure is so informal” (I5) and, ultimately, “power in Wikipedia is: sufficient people listen to you and are inclined to consider what you want done” (I10). Boundaries that define community relationships are flexible and indistinct; however, Wales suggested that community boundaries are undergoing a process of demarcation (Ostrom Principle 5):

One of the things I think we’re beginning to see is a much clearer sense of self-identity of the core community versus the broad community. And some scuffles around certain topics where the core community feels like there are things in Wikipedia that should not be there or practices that are engaged in that should not be engaged in.

Registered users often self-select into formal and informal subgroups along ideological, functional, and content-related lines. Ideological groups are much like political parties whose affiliates hold a set of common beliefs about the way the community should function and what its goals should be. Examples of stable ideological groups include deletionists, who are committed to very strict guidelines on what constitutes encyclopedic topics, and inclusionists, who are committed to the idea that an online encyclopedia need not and must not exclude information. Functional groups can be thought of as wiki professions. Examples of functional groups include vandal-fighters, who spend tremendous amounts of time reverting vandalism and managing vandals, committees such as the Bot Approvals Group, who review and approve software that can automatically edit the site, and Open Ticket Request System (OTRS) volunteers, who manage the email help-ticket system for requests from the public. Stable content-related subgroups are coalitions of editors with common interests or expertise; these have proliferated widely as “WikiProjects” and will be examined in detail in later sections.

Registered users may also hold various technical powers that are generally understood to be distinct from the social influence they exert as individual community members. The main technical distinctions among users in English Wikipedia include administrator, bureaucrat, checker, oversight, developer and steward (See Table 1). Interviewees explained that:

A lot of people look at the different levels of power in the software and assume that those also correlate to levels of power in the community, when in fact that’s generally not true... Those powers are more or less independent of

community powers. They are considered technical functions, they are considered janitorial work. (I11)

It's supposed to be a technical—above anything else—distinction. (I7)

We like people who have those powers to be trusted by the community. But it's not really supposed to be seen as a status position. Because we don't really want to have classes of people. So it's more of people who are trusted but not really somebody's status per se. (I4)

In practice, we found that technical and social power cannot be entirely uncoupled. In fact, people who retain the most powerful technical capabilities on the site are always individuals who are highly trusted and well known by established members of the community. As such, they tend to be capable of exerting considerable influence socially. The main danger in equating technical and social power is that it leads one to overlook the potentially enormous power that can be held by regular users of Wikipedia.

Table 1: Access Levels in English Wikipedia

Administrator	Protect/unprotect pages; Delete/undelete pages; Block/unblock users; Special revert tools
Bureaucrat	Make administrators; Rename users; Make other bureaucrats
Steward	Change all user access levels on all Wikimedia projects
Oversight	Hide page revisions from all other user types
Checkuser	View user IP addresses
Developer	Access to MediaWiki software and Foundation servers (various sublevels)

The *Arbitration Committee* wields considerable influence in the community. The Arbitration Committee (Arb Com) was conceived of as the last step in a formal dispute resolution process put into place in early 2004; however, today it appears to often serve as a more general decision-making body for the English language site. Arb Com was initially charged with interpreting policy and making binding resolutions in the case of interpersonal disputes. Explained one Arb Com member: “It’s kind of a quasi-legal type of thing. Some people like it, some people don’t, but hey, it works” (I7). Committee members are selected through a hybrid process of election by the community and appointment by Jimmy Wales. Arbitrators have no special authority with respect to content or any formal power to create policy, yet we will see that Committee action can play a role in influencing both policy and content.

Jimmy Wales plays a leadership role in the community that has also changed over time. In the past he was more involved with the daily business of the site.

His current involvement in English Wikipedia is mostly limited to high-level guidance and policy discussions.

Finally, English language Wikipedia is one of several hundred wiki projects served by the Wikimedia Foundation, which owns the servers that run it and manages financial assets. The Foundation does not directly make content decisions about the various language projects or manage local governance decisions; its main influence on English Wikipedia is in setting high level goals for and sometimes legal constraints on the projects, which are implemented locally as the community sees fit.

Table 2: Selected Interviewee Roles

Site and Wikimedia Foundation founder, Jimmy Wales
Arb Com member/former member (5 interviews)
Regular user (1 interview)
Involved in Wikimedia Foundation (3 interviews)
Volunteer contributor to MediaWiki software (1 interview)
WikiProject founder and/or participant (6 interviews)
Users with combinations of access levels including: administrator, checkuser, oversight, bureaucrat, steward, developer (10 interviews)

4. Policy in Wikipedia

Wiki policy is fluid. Traditionally, it has tended to echo community practices rather than prescribe them. Explained one Arb Com member, “even though they’re written down, anybody can edit a policy page even, because it’s a wiki. That includes every single policy that exists. So it’s possible on any given day that a policy—this is what’s kind of odd, I think—can not really reflect exactly what is right” (I4). This openness might initially give the appearance that policy on Wikipedia is “up for grabs” by anyone who wanders along when, in fact, the creation and refinement of policy is a complex social negotiation that often takes place across many communication channels and in which power, authority and reputation play decisive roles. We found that on-wiki activity only accounts for a portion of governance activity on the site. Often, critical decisions are discussed in public and private IRC channels, mailing lists, personal email, and other off-wiki communication.

Policies are not the only “legal” artifacts on the site. *Guidelines* are strong recommendations for behavior, content, stylistic conventions and the like that are not followed as strictly as policy. Policies and guidelines are similar in that they are established by consensus and require broad community support. Other artifacts of governance include essays, policy proposals, rejected proposals, Arb Com decisions, and pages of historical significance. In several dimensions, we found that governance mechanisms around policies and guidelines

are becoming increasingly decentralized as the site becomes larger (see Figure 1).

	of content-related policy/guidelines	of behavior-related policy/guidelines
Creation		has slowed
Modification	has become increasingly decentralized (WikiProjects)	is done by consensus
Interpretation		has become increasingly decentralized (Wales -> ArbCom -> Admins)
Enforcement	is done by consensus (WikiProjects have no local authority)	remains distributed among admins

Figure 1: Trends in Management of Wikipedia Policies and Guidelines

4.1 How is Policy Made and Changed?

There is no one way Wikipedia policy gets made. Wikipedia policy making was not part of a planned system of governance; each policy arose from a unique set of circumstances and pressures from different stakeholders and regulating influences at different points in the site’s history. Several of the longest standing Wikipedians that we interviewed indicated that the site has undergone deep changes with respect to the creation and role of policy. Here we will describe these changes and the patterns in our interviewees’ stories that point to consistent governance mechanisms in the site. We found that site-wide policy making has slowed and mechanisms that support that creation and improvement of guidelines—in particular, stylistic and content-related conventions—have become increasingly decentralized.

We opened the paper with a question about whether or not Wikipedia should include an article for each victim of the Virginia Tech shootings. Before we discuss governance mechanisms and trends, we will give a concrete example of policy creation and evolution that impacts decision making about biographical material.

4.1.1 Case: *Biographies of Living Persons*

How Wikipedians write biographical material has been an object of continuous reform. The story of the policy on Biographies of Living Persons (BLP) is interesting because it illustrates the thick tangle of circumstance that gives shape to policy on the site. Different Wikipedians give different weight to the factors that played a role in its formation, but by 2005, growth in the site’s popularity had set the stage for several situations that influenced adoption of the policy.

One development was that people began to notice and edit their own biographies, which led to conflicts with Wikipedians and a flurry of complaints to the Wikipedia help system (OTRS):

Really how the Biography of Living Persons policy was developed was a lot of conversation on the OTRS mailing list. We had all these people who were complaining that their biography was all messed up. They tried to make some changes to their biography and they kept being reverted for various reasons, whether they were right or they were wrong. (17)

Another development was that Wikipedians themselves began to feel a sense of responsibility for the quality of information they were offering about people.

This is another reason why Biographies of Living Persons was created, because we had so many problems with all these unsourced statements and people used to only put like, at most they’d mention on the talk page “ok we need a cite for this particular issue” or they put “citation needed” in the text itself. Now we have the policy that we have to immediately get rid of that. (17)

To complicate matters, individual Wikipedians were applying social norms from their offline social worlds to make judgments about what was acceptable editing behavior, which led to internal conflicts:

I was concerned about the way that people were being written about in the encyclopedia... And so I was looking through and noticing that people were being written about in ways that I thought were not correct. So I was removing some people out of categories that I didn’t feel they belong in that were not supportive and, you know, that type of thing... And I got into a disagreement with a couple of editors over this. And it turned into an arbitration case. (14)

Finally, the most publicly evident situation was pressure mounting on the community from external sources. For a non-profit organization dependent on the charitable contributions and good-will of the public, bad press amounts to a hard market constraint on behavior:

There have been a series of policies on how we handle biographical articles and unsourced assertions that have come about from a series of negative incidents in the press involving Wikipedia and erroneous information in biographies. (15)

Eventually, after much off- and on-wiki discussion about the situation, a proposal page was started and the community began constructing what was initially a proposed guideline. Eventually, the page reached a form acceptable to most community members. The guideline became strengthened in part because of the arbitration case mentioned above, although the move to policy was also discussed on the policy talk page:

The way that the policy went from being a guideline to a policy was it was cited by arbitration, it was used by the community. Originally the Biographies of Living Persons was just a guideline. (14)

After a year and a half, this topic continues to instigate policy reform on the site. There is steady discussion of BLP-related issues on public mailing lists and, recently, a case was accepted by Arb Com that concerned the creation of articles about private citizens:

People have begun to understand more and more how we affect people's lives. Particularly people that are not really very public people who just, you know, by happenstance are famous for a little while for maybe some event in their life. (I4)

As the sheer size of Wikipedia has gotten larger and larger, this means that we're now writing articles on less and less famous people and that can be a lot more personally distressing to people. (I11)

The day before he was interviewed for this study, Jimmy Wales made a policy change that affected biography-writing practices on the site. Wales amended the policy "What Wikipedia is Not" to specify that Wikipedia is not a newspaper. He explained:

I added Wikipedia is not a newspaper and especially not a tabloid newspaper and that we... attempt to make some sort of judgment about the long term historical notability of something... And that was just a change that I made yesterday just by boldly updating the policy page. And will that stick or not?

Weeks later, the policy has been refined, copyedited, and extended to include heuristics for determining long-term notability. Its essential message remains unchanged. The policy stuck—for now.

This story illustrates the composite and evolutionary nature of policy development in Wikipedia, and the varied roles of regular users, Arb Com, Jimmy Wales, and external forces. In addition to illustrating the emergence of policy, our interviewees' comments point to emergent social norms on Wikipedia that compel users to write with respect for the privacy and well-being of human beings who are the subject of encyclopedic material. Although they are no longer living, these same issues have influenced decisions about biographies of Virginia Tech victims. After a five-day on-wiki discussion consisting of roughly 18,000 words and more than 400 edits, a proposal to delete a list of victims was rejected; the list is currently maintained with minimal biographical detail. It is felt that although the event is noteworthy, including details of individual victims' lives is gratuitous—biographies that focus on the circumstances of one's death are not balanced or encyclopedic. Consensus emerged that these biographies are subject to the same standards of notability as other biographies on the site. In the next sections, we will inspect the mechanisms by which policy is created, changed, interpreted and enforced on the site.

4.1.2 Policy-Making Mechanisms

In the early days of Wikipedia, the community was small and policy making was less formal. One interviewee mentioned that the editor-in-chief of Nupedia, Larry Sanger, played a large role in guiding the earliest Wikipedia policies. Because we did not interview many individuals who were active in that stage of the project and because we do not have access to the earliest policy-related wiki pages, we do not have detailed data about those earliest policy making efforts. We do know that early policy making on the site was comparatively informal and done in an ad-hoc fashion. This process has become more formalized over time as the site grew and informal consensus became harder to manage.

In the early stages it was just a matter of just a couple people getting together and saying "Does this work? No not really. Does this work? Yeah, I think that's good... Let's just do this. At least for now." And that "at least for now" part is really ironic because now it's hard set policy. (I7)

When we made that [undeletion] policy there wasn't a policy on how to create policy, which there is now. [laughs] So it was done very informally, just people discussing it together... The reason [the policy on how to create policy] had to come about is that people just kept writing pages up and then other people didn't know whether it was policy or not. And in the early days that just wasn't a problem because you could just go in and discuss, ok, is this really policy or is this just someone's idea? But as the project got larger, new people coming in and reading these pages and not knowing whether it was just someone's random thoughts one day or whether the whole community agreed with it. So they got a little bit more formal about tagging things as policy or policy-in-development or rejected policy or whether it's just an essay and not a policy at all. (I1)

Wales noted that today there are a few common forms of policy making:

You can have a community-wide vote with a fairly overwhelming majority and that will normally then cause something to become policy. You can have someone just boldly going in and changing policy and if it sticks, it sticks... The last way that policy can get made is I just say so. And that's done very rarely and I don't do it unless I feel I have strong support from the community.

Although he has introduced policy and policy modifications, Wales' changes have generally already been discussed by many individuals. His support of a policy appears to be more often a culmination than a germination of policy discussion.

Some interviewees suggested that policy making has not just become more formal, but that policy making efforts have slowed in recent years:

I think that as the site got older and there was a lot of policy there already it wasn't necessary to create new policies even though people are still doing that. (I1)

The reality is there's been very little change in policy since Wikipedia became wildly popular. If you look at where policy was at in say the beginning of 2005 and where it's at today there's very little difference. What's happened is that since policy making is by consensus and the number of active contributors has grown into the 1000s, it's almost impossible to achieve consensus anymore. (I5)

The process for removing policy is embedded in the detailed specifications for removing articles from the site. A close examination of this process is beyond the scope of this paper; however, it is important to note that article deletion is usually grounded in the "WikiProcess" of discussion and consensus building, with some exceptions for blatant examples of vandalism or content-related policy violations. Consensus building among interested parties is the foundation of nearly all decision making on Wikipedia (Ostrom Principle 2).

4.1.3. Decentralization in Policy Creation

The problem of achieving consensus about content guidelines as the organization grows has led to a proliferation of small, decentralized social structures. These nodes, called WikiProjects, are particularly important when it comes to developing guidelines for stylistic conventions and the creation of content. One might think of these as local jurisdictions in the site, within which local leadership, norms, and standards for writing are agreed upon by editors familiar with a particular topic (Ostrom Principle 1). WikiProjects are most commonly organized around subject matter; however, many WikiProjects for specific tasks such as creating audio versions of articles (Spoken Wikipedia), or identifying and eliminating bias (WikiProject Countering systemic bias) also exist.

One thing that is worth mentioning is that there's a lot of niche areas in Wikipedia... Often times it comes down to what do the people who are interested in that particular niche, what is the standard way of doing it will come down to what do those 4 or 8 people agree on. (I6)

The fact that it's gotten so big, it's hard—people can't keep up to date with everything that happens anymore. So in a sense I think the growth of WikiProjects has been partly a reaction to that. In that, you know, you no longer feel part of a community of thousands and thousands of people. So people set up smaller communities. (I10)

As it gets larger it's kind of hard to know everyone who's there, so... people try and find new ways of dealing with issues when they don't know everything that's going on... from the content side, WikiProjects are formed to kind of focus on articles in a particular area. And they can develop policies that just relate to their area. So it enables people to still work together within a small community

even though Wikipedia itself is a really huge community. So I think this kind of breaking up within Wikipedia itself is one of the ways that it's changing. People are organizing themselves into smaller groups so they can still work the same way they did originally. (I1)

WikiProjects run into governance issues themselves when two or more projects claim that a particular article falls within their province or when there is a need to enforce local policy. Essentially, no governance mechanisms exist to address either of these situations directly and users employ the standard wiki process of discussion and consensus building. When it comes to local disputes, keeping up good relations with neighboring projects is important. One interviewee explained that he acts as the informal organizer for one WikiProject and that this involves liaising with projects in related subject areas "to make sure that everybody is on the right track when editors from other projects come and edit [our project] pages, this kind of thing." (I2)

Another interviewee explained that not long after starting a WikiProject, he left because participants were not interested in following guidelines that he felt ensured quality content production. He told us that:

In the case of general guidelines or policies that exist for Wikipedia articles under normal circumstances, people can get chastised or ostracized by the community for constantly breaking those ... Now in the case of the [WikiProject] obviously you can't do that" (I8)

Another WikiProject founder observed that WikiProject guidelines are nested within but must not deviate from general Wikipedia guidelines: "we always have to base our guidelines on the ones of Wikipedia. And of course, we can't become a unique wiki inside the whole Wikipedia" (I9) (Ostrom Principle 8). One interviewee suggested that the authority of WikiProjects to manage themselves is viewed with some misgivings:

I suspect that as the site gets bigger, WikiProjects will have more of an influence because there can't be as much central authority, the place is simply too big. Groupings of editors interested in a particular topic will—that may also be a problem in that groups of editors might decide to overwrite important site-wide goals. And I think that will be a source of conflict in the future. (I10)

So, for the time being, WikiProjects exist as an administrative structure that carries out many of the functions related to organizing content production and setting editorial standards and stylistic convention; however, they lack the ability to compel editors to adhere to local rules (Ostrom Principle 7).

A counterexample to the trend of decentralization in Wikipedia can be found in the case of policies regarding behavior on the site that could result in harm to individuals or represent a direct legal or financial threat to the Wikimedia Foundation. Some such policies are simply "inherited" by the Wikipedia project from the Foundation, such as rules regarding copyright. Sensitive

issues regarding children or illegal activity have also motivated unilateral decision making by Arb Com and Jimmy Wales. Decisions that threaten the well-being of project contributors or the health of the Foundation may result in policy setting that is uncharacteristically opaque, centralized and not open to discussion.

Policy making has in one sense always been highly decentralized because it requires consensus of the community to stick. As we will see in the next section, even if some policies are initiated as unilateral decisions by central authority figures, it requires the consent of the governed to actually become policy because enforcement remains highly decentralized.

4.2 How is Policy Interpreted and Enforced?

We began the previous section with the story of policy that regulates biographical content production on Wikipedia. In order for BLP or any other Wikipedia policy to matter, someone needs to decide when it applies and enforce it. When it comes to interpretation and enforcement, Wikipedians make some distinctions between content-related and behavior-related policy. Interpretation of content-related policy tends to be highly decentralized. The entire community monitors content and if a dispute over arises, it is generally decided through discussion and consensus building by the people involved in the situation (Ostrom Principle 3). When behavior-related policy is broken, a series of graduated sanctions can be imposed that begin with the posting of warnings and can lead to banning from the site (Ostrom Principle 6). If a dispute or behavior-related policy infraction cannot be resolved locally using mild sanctions by the people directly involved, it can be referred to a formal, centralized dispute resolution process with the authority to impose more severe punishments. We found that growth of the community has resulted in increasing decentralization of the final stages of dispute resolution and use of severe sanctions.

4.2.1 Decentralization in Policy Interpretation

When the community was young, Jimmy Wales often handled interpersonal disputes by fiat, so authority was as centralized as it possibly could be. As the community grew, his ability to invest the time necessary to make good decisions was reduced. Eventually it was necessary to create a formal process and a kind of judicial body that could interpret official policy in the best interests of the project.

Originally Jimmy Wales had the power that the Arbitration Committee has. You know, the power to ban people, the power to put restrictions on people above and beyond the ordinary ones imposed on everybody. So in a sense our power is the devolved power from the king so to say. (I10)

There was no way that Jimmy was able to keep up with all the type of crap that was going on really and he needed help. (I7)

As with the reification of practice in other policies, the formalization of dispute resolution also alleviated the difficulty of communicating common practices to newcomers, which was important for the growing community.

As the site got larger... people who were having a problem didn't really know who to go to. It's not like when there's a small community, you know who to go to because you know who the lead people in the community are. Once it gets larger and you don't know everyone anymore, you don't really know if there's anyone to turn to, what to do about this problem that you're having. So I think it made it a lot easier for new people who are coming in and finding themselves in some sort of fight to get help from a more experienced editor. (I1)

We understand from interviewees that Arb Com procedures underwent an initial period of refinement, but it seems that after nearly three years of operation, Committee procedures have stabilized as a quasi-legal process of collecting evidence and statements, deliberation, and issuing findings of fact, principles, and remedies.

The role of Arb Com seems to have changed from a dispute resolution body to a more general decision-making body and counsel. Although arbitrators were not charged with creating policy, and cannot officially take policy-making action, their actions may have far-reaching repercussions on how policy is interpreted and even on the creation of policy in the site.

So with the creation of the dispute resolution process, final interpretation of behavior-related policy moved out from Wales to Arb Com. Today it is in the midst of further decentralization. Despite the traditional division between technical and social powers on the site, administrators are beginning to step into more authoritative roles and are making more and more interpretive and "moral" decisions about user behavior. Nearly every interviewee suggested that, for better or for worse, the role of administrator carries with it more social authority today than it ever has in the past.

Originally the whole idea of administrator was seen as somebody that was seen as more custodian, you know? And they have in some ways turned into more—I mean there's a group of administrators who work at more policing problems, disruptive problem editors and don't just do like cleaning up vandalism and removing stuff that is just not appropriate. (I4)

As time goes on, we're now to the point where the arb com can't possibly hear all of the cases that come up and so you see administrators getting together having ad hoc decisions about what to do with a user and kicking people off of the site when they think it's justified and when they think they'll have support from the arbitration committee and the community at large. -(I5)

Arb Com is... taking fewer less complex cases and leaving the easier ones to sort out to the administrators. That has been the trend. (I6)

We have the development that the administrators more and more are assuming the roles and the discussing the things and making the decisions. Because we're just a small group of people and we hear a few cases but the bulk of the action is actually in the larger committee of the administrators. (I3)

It's less and less of a janitorial role and it's becoming a little bit more, at least from what I've seen, more and more of an independent—especially since the administrator's notice board was created. A bunch of administrators get together and they decide on particular things. They're kind of deciding more and more on the type of things that the early Arb Com used to do on how to handle certain cases. (I7)

One might argue that the cost associated with dispute resolution (Ostrom Principle 4) is further reduced as it is further devolved among members of the community.

4.2.2 Enforcement

The blurring of the distinction between social and technical powers of administrators means that a substantial amount of power is consolidating in one section of the Wikipedia community. Administrators, after all, are the traditional enforcers of policy. As of June 2007, there are approximately 1200 administrators in the English Wikipedia. In the past they relied on community consensus or Arb Com to make decisions about which users should be blocked, which pages should be protected, and which pages should not exist. If administrators are stepping into the role of interpreter of policy, they are positioned to wield what some consider excessive authority over Wikipedian behavior. The corollary to this development is that requirements to become an administrator on the English language site have become increasingly rigorous.

They go through this ridiculously insane process in order to become an administrator... I don't even know if I could become an administrator now. (I7)

Nowadays it is much harder to become an administrator because the standards have risen... the consensus seems to have emerged that, firstly, once someone is admin they should not be revoked again, and secondly, as a corollary of that we should very carefully select our admins. (I8)

The distributed model of community enforcement has traditionally provided a check on the power of any governing body in the site. Arb Com, Jimbo, or any other segment of the community must act with widespread support or decisions simply do not get enforced.

We post our decision and let the administrators go ahead and enforce it. (I6)

The main limit on the power of the arbitration committee is the fact that we rely on the community for enforcement... I think almost every major decision that

we've encountered there has been—the committee discussed one or another remedy that has had to be abandoned because we didn't think the community would go along with it. (I5)

Since it's generally accepted as a policy that admins don't have to do anything, the fact is, if we made an unpopular ruling the admins can simply vote with their feet and refuse to enforce it. (I10)

A noteworthy example of community solidarity in declining to enforce an Arb Com decision can be found in the case of William Connolley. William Connolley is a respected British climatologist who makes many contributions to the encyclopedia in the area of global climate change. One Arb Com member pointed out that “He has had to deal with some really problem users. You can imagine if you edit on the global warming article the kind of crap that comes his way.” (I6) Connolley eventually got into a dispute with another editor and the case was accepted by Arb Com. The case was problematic because Connolley had broken behavior-related policies in order to defend the content of the encyclopedia. One Committee member explained how the ruling was made to avoid the appearance of Arb Com involvement in content decisions:

William Connolley is citing established journals for specific facts. This other person is citing Michael Crichton and periodicals and crackpot webpages. We had to make a distinction—we had to use the existing policy on references to say ok, “this other person who is citing Michael Crichton is not following this policy to use science-related references which have to be to established journals but William Connolley is. Therefore the behavior of this other person on citing these poor sources is going to be penalized” We weren't saying “Connolley is right, this other person is wrong, this other person needs to get hit on the head.” (I7)

But Connolley had broken rules as well. Explained another Committee member, Connolley “doesn't suffer fools. Well, so we said ‘look, we've got lots of fools here. You're just going to have to suffer fools. [laughs] You're going to have to be polite to them anyway.’” (I3) The decision was made that Connolley was restricted to one revert per day, which essentially meant that he could only remove another user's edits once per day. A third Arb Com member explained this action had the effect of “restricting his ability to deal with those problem users.” (I6) Many members of the community disagreed that these restrictions should have been placed on Connolley; as a result, the Arb Com decision was not enforced. “He just raised hell and everybody else raised hell and finally we had to say, well, that was a bad idea.” (I3) Arb Com eventually reversed its decision, but not until the term of parole had nearly expired. Still, the action is emblematic of the power held by administrators.

The administrators who should have been blocking him if he did that or whatever, said “I don't want to block him. I

don't feel like it's harmful..." Administrators are not like police officers. You're not sworn to block people, you're not sworn to carry out things. Nobody is. So if no one does it then it's not something that can be enforced. But more or less in this particular case, that is what happened. And eventually it got lifted because it was said "nobody wants to do this and no-one thinks it was a good decision." And so it ended up being reversed. (14)

5. Conclusions

Governance is not simply the sum of policy-related social processes. There are also groups of individuals who manage the execution of detailed procedures that operate on a site-wide level such as the Featured Article (FA) selection process used by Viegas et. al. to reveal features of Ostrom's eight principles of self-organization in Wikipedia. The story we have told about policy in Wikipedia provides a complementary view. The most noticeable gap we found between Ostrom's principles and the structure of Wikipedia governance was in the inability of WikiProject members to enforce their own local guidelines in order to maintain local standards of content production. It remains to be seen whether this trend will prove disruptive over time. Perhaps WikiProjects or some other form of local governance in the site will acquire the authority to impose sanctions on users who violate their locally-devised rules. It is also possible that local enforcement will simply prove unnecessary and that Wikipedia will ultimately demonstrate an alternate model of self-governance for which theoretical justification will need to be found. Rather than suggest a prescribed community structure based on discrepancy with existing theory, we can continue to learn from the evolving practices of Wikipedians.

Issues like the Virginia Tech biographies and arachnophobia illustration arise every day, and Wikipedians continue to use and refine both their policies *and* their policy-related processes. Overall, the story of Wikipedia governance that we assembled is one of increasing decentralization. As the community grows, it has become necessary for governance mechanisms to shift outward into the community. This decentralization was not entirely accidental; self-organization was dependent in part on the design of the technology and embedded in the philosophy of the community's founder and early participants.

We suggest that the Wikipedia community has remained healthy in large part due to the continued presence of "old-timers" who carry a set of social norms and organizational ideals with them into every WikiProject, committee, and local process in which they take part. Instead of fracturing, the community has (so far) gracefully distributed the pressure of expansion among its members in ways that are largely consistent

with Ostrom's propositions about the necessity of decentralized decision-making authority in large, self-organizing enterprises.

6. References

- [1] Benkler, Y. Coase's Penguin, or, Linux and The Nature of the Firm. *The Yale Law Journal*. vol.112, no.3, 2002, 369-446.
- [2] Benkler, Y. *The Wealth of Networks*. Yale University Press, London, 2006.
- [3] Bruckman, A. Studying the amateur artist: a perspective on disguising data collected in human subjects research on the internet. *Ethics and Information Technology*. vol. 4, no. 3, 2002, 217-231.
- [4] Castronova, E. *Synthetic Worlds. The Business and Culture of Online Games*. University of Chicago Press, Chicago, 2005.
- [5] Cosley, D., Frankowski, D., Kiesler, S., Terveen, L. and Riedl, J. How oversight improves member-maintained communities. *CHI* (Portland, OR) 11-20, 2005.
- [6] Dibbell, J. *My Tiny Life: Crime and Passion in a Virtual World*. Henry Holt and Company, New York, 1998.
- [7] Lessig, L. *Code 2.0*. Basic Books, New York, 2006.
- [8] Morningstar, C. and Farmer, F. The lessons of Lucasfilm's habitat. *First Annual International Conference on Cyberspace*, 1990.
- [9] Ostrom, E. Collective action and the evolution of the social norms. *Journal of Economic Perspectives*. vol. 14, no. 3, 2000, 137-158.
- [10] Ostrom, E. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, Cambridge, 1990.
- [11] Preece, J. Etiquette, empathy and trust in communities of practice: stepping-stones to social capital. *Journal of Universal Computer Science* vol. 10, no. 3, 2004, 194-202.
- [12] Schutz, A. *The Phenomenology of the Social World*. Northwestern University Press, Evanston, IL 1967.
- [13] Seidman, I. *Interviewing as Qualitative Research*. Teacher's College Press, New York, 1991.
- [14] Smith, M. and Kollock, P. *Communities in Cyberspace*. Routledge, London, 1999.
- [15] Sunstein, C. *Infotopia: How Many Minds Produce Knowledge*. Oxford University Press, Oxford, 2006.
- [16] Viegas, F., Wattenberg, M., Kriss, J. and van Ham, F. Talk before you type: coordination in Wikipedia. *Hawai'ian International Conference on System Sciences*, 2007.
- [17] Viegas, F., Wattenberg, M. and Dave, K. Studying cooperation and conflict between authors with history flow visualizations. *CHI* (Vienna, Austria), 575-582, 2004.
- [18] Viegas, F., Wattenberg, M. and McKeon, M. The hidden order of Wikipedia. *HCI* (Beijing), 2007.
- [19] Weber, M. *Basic Concepts in Sociology*. Citadel Press, New York, 1966.
- [20] Winner, L. Do artifacts have politics? *The Whale and the Reactor: A Search for Limits in an Age of High Technology*, 1986, 19-39.