Pages 1 - 19.

OPEN ACCESS

Article

The Evolution of Wikipedia's Norm Network

Bradi Heaberlin[1,2], Simon DeDeo[1,3,4,5]*

[1] Program in Cognitive Science, Indiana University, 1900 E 10th St, Bloomington, IN 47406, USA
[2] Department of Political Science, Indiana University, 1100 E 7th St, Bloomington, IN 47405, USA
[3] Center for Complex Networks and Systems Research, Department of Informatics, Indiana University, 919 E 10th St, Bloomington, IN 47408, USA
[4] Ostrom Workshop in Political Theory and Policy Analysis, 513 N Park Avenue, Bloomington, IN 47408, USA
[5] Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA

* Author to whom correspondence should be addressed; sdedeo@indiana.edu, telephone (812) 856-2855, fax (812) 856-3825

Version December 8, 2015. Typeset by LaTeX using class file mdpi.cls

---

Abstract: Social norms have traditionally been difficult to quantify. In any particular society, their sheer number and complex interdependencies often limit a system-level analysis. Here, we present the development of the network of norms that sustain the online social system Wikipedia. We do so using a network of pages that establish, describe, and interpret the society's norms. Fifteen years of high-resolution data allow us to study how this network evolves over time. Despite Wikipedia's reputation for ad hoc governance, we find that its normative evolution is highly conservative. The earliest users create norms that dominate the network and persist over time. These core norms govern both content and interpersonal interactions, using abstract principles such as neutrality, verifiability, and assume good faith. As the network grows, norm neighborhoods decouple topologically from each other, while increasing in semantic coherence. Taken together, these results suggest that the evolution of Wikipedia's norm network is akin to bureaucratic systems that predate the information age.

Keywords: social norms; norm networks; Wikipedia; oligarchy; bureaucracy; governance; knowledge commons

---

1. Introduction

A society's shared ideas about how one "ought" to behave govern essential features of economic and political life [1–6]. Outside of idealized game-theoretic environments, economic incentives are supplemented with norms about honesty: a higher wage is possible when workers believe they ought not to cheat their employer [7]. States can make laws, but people determine their legitimacy based on beliefs about fairness and authority. A police force without legitimacy cannot enforce the law [8,9].

Norms are under continuous development. The modern norm against physical violence, for example, has unexpected roots and continues to evolve [10–12]. Yet we understand far less about the history and development of norms than we do about economics or the law [13]. We often lack the data that would allow us to track the coevolution of complex, interrelated, and interpretive ideas such as honesty, fairness, and authority the way we can track prices and monetary flows, or the creation and enforcement of statutes.

Online societies such as Wikipedia provide new opportunities to study the development of norms over time. Wikipedia's minute-by-minute server logs cover nearly fifteen years of social evolution in a population that, for a large fraction of its history, has numbered in the tens of thousands. We focus on a special subspace of the encyclopedia devoted to information and discussion about the project itself. This allows us to quantify how editors describe expectations for behavior, and in doing so both create and reinterpret the norms of their society.

Along with information and code repositories at the center of the modern global economy, such as GNU/Linux, Wikipedia is a canonical example of a knowledge commons [14–17]. Knowledge commons rely on norms, rather than markets or laws, for the majority of their governance [18,19]. Norms matter on Wikipedia in ways that make it impossible for participants to ignore. It is norms, rather than laws, that dictate what is or is not included, who participates, and what they do.

We study the evolution of norms on Wikipedia using a special subset of tightly-linked pages that establish, describe, and interpret them. Instead of a rigid list of regulations which users are expected to uphold without question, these pages form a complex network of policies and commentary under continuous development since the encyclopedia was created. Our network perspective allows us to go beyond the tracking of a single behavior over time (a common approach in studies of cultural evolution [20]) to look at the evolution of relationships between hundreds, and even thousands, of distinct ideas.

We use this data to ask three critical questions. In a system where norms are constantly being discussed and created, how, and when, do some norms come to dominate over others? What kinds of behavior do they govern? And how do those core norms evolve over time?

Our answers are surprising. While some accounts of Wikipedia stressed its flexibility and the ad hoc nature of its governance [21–23], we find that Wikipedia's normative evolution is highly conservative. Norms that dominate the system in Wikipedia's later years were created very early, when the population was much smaller. These core norms tell editors both how to write and format articles and how to collaborate with others when faced with disagreements and even heated arguments. They do so by reference to universal, rationalized principles such as neutrality, verifiability, civility, and consensus. Over time, the network neighborhoods of these norms decouple topologically. As they do so, their internal semantic coherence rises, as measured using a topic model of the page text. Wikipedia's abstract core norms and decoupling process show that it adopts an "institutionalized organization" structure akin to bureaucratic systems that predate the information age [24].

2. Methods

To gather data on the network of norms on Wikipedia, we spider links within the "namespace" reserved for (among other things) policies, guidelines, processes, and discussion. These pages can be identified because they carry the special prefix "Wikipedia:" or "WP:". Network nodes are pages; directed edges between pages occur when one page links to another via a hyperlink that meets certain criteria. We begin our spidering at the (arbitrarily selected) norm page "Assume good faith". Editors can classify pages by adding tags; these tags include, most notably, "policy", "guideline", and "essay", among others. When we download page text, we also record these categorizations. Details of this process, our filters, and our post-processing of links between pages, appear in the Appendix.

Based on, and extending, pre-existing literature on the nature of norms in the knowledge commons [19], we consider three distinct norm categories. Norms may attempt to regulate content creation ("user-content" norms) and interactions between users ("user-user" norms). In addition, norms may attempt to define a more formal administrative structure with distinct roles, duties, and expectations for admins ("user-admin" norms). The two authors of this paper independently categorized a random sample of forty pages using this scheme, and we calculated inter-coder reliability using Cohen's kappa [25].

For our semantic analysis, we include all text except that found in special boxes whose text is replicated by template across multiple pages. To build our distribution over one-grams, we normalize all text to lowercase, merge hyphenated words ("error-correction" to "errorcorrection"), and drop punctuation. We do neither stemming nor spelling correction.

A critical external variable is the number of active users on the encyclopedia at any point in time. Following Ref. [26], we define an active user as one who has made five or more edits within a month; these statistics are publicly maintained at Ref. [27].

2.1. Attention Measures

The pages in our corpus are created with the explicit goal of explaining the norms of Wikipedia to editors, and of influencing their interactions with others. Users navigate the system as a network structure, and consequently encounter some pages more than others.

We measure this using Eigenvector centrality (EC), which quantifies the importance of a page based on its overall accessibility within the network. The EC of a page is the probability of happening across a page during a random walk; equivalent to the PageRank algorithm, it is used in the behavioral sciences to identify consensus on dominance and power [28]. We set ϵ, the probability of a random jump, to 0.15.

We expect some pages to attract high levels of attention while others are largely peripheral to the network. We quantify this inequality using the Gini Coefficient (GC). GC varies between zero (perfect equality; all pages have equal EC) and one (one page has a high EC, all other pages have the same low value). GC is widely used in economics to measure income inequality. Here, it provides a global measure of the extent to which a system is dominated by a few norms. As a dimensionless quantity, it allows researchers to compare this system to others that might be the subject of later research.

2.2. Influence and Overlap

Editors provide a context for each page in the network by the links they make to other pages. A node p can be understood to *influence* a node q when q links to p. To quantify the influence of page i, we consider the nodes a random walker will encounter when beginning at i and following the links in-bound to her current node.

More formally, placing a random-walker at node p we allow her to take n steps from this starting point along the direction-reversed network; the probability distribution over the network we can write as p_i, the probability of the walker ending up at node i. A distribution p_i defines the influence of node i.

To quantify the distance between two nodes, we then consider the influence *overlap* between two arbitrary nodes p and q. Overlap quantifies the extent to which two random walkers, beginning at these nodes, will tend to visit the same pages. If p_i and q_i are the probability distributions associated with the influence of node p and q, then overlap is defined as

$$O(p,q) = \frac{\sum_{i=1}^{N} p_i q_i}{\left(\left[\sum_{i=1}^{N} p_i^2\right]\left[\sum_{i=1}^{N} q_i^2\right]\right)^{1/2}}. \quad (1)$$

For multiple pages, we can compute the average pairwise overlap simply by averaging the overlap between all possible pairs within the set.

High overlap between p and q indicates that two pages influence a large number of common nodes. When n goes to infinity, the random walkers converge to the stationary distribution and overlap is one; conversely, when n is small, random walkers have less time to encounter each other. We take n equal to five—larger than the average shortest path (roughly three, in our network), so that nodes are potentially reachable, but much less than the convergence time to the stationary distribution.

Overlap can be thought of as a measure of distance. It invokes only local mechanisms: users traveling from one page to another by the links that connect them. This is in contrast to a measure such as shortest paths, which is computationally expensive and requires detailed, global knowledge of the network link-structure. In general, for example, the number of nodes an algorithm needs to visit in order to determine the shortest path between two nodes will usually be much larger than the length of the final path.

Both influence and overlap require us to specify particular nodes of interest; we focus in this work on pairs of high-EC pages, or core norms.

2.3. Semantic Coherence

We consider the semantic relationships between pages. This provides a notion of relatedness distinct from how norms connect via hyperlink relations. To do this, we do topic-modeling (Latent Dirichlet Allocation [29]) on the one-grams of the visible, human-readable text on each page. Topic models allow us to represent short texts even when they draw from a rich vocabulary: topics coarse-grain the underlying distributions over words. Underlying parameters of the model are reported in Appendix Sec. E.

With the resulting topic model, we can then compute the semantic distance between all pairs of pages using Jensen-Shannon Distance (JSD), a measure that quantifies the distinguishability of two distributions [30]. This gives us a weighted semantic network that we can compare to the network of hyperlinks between pages. In particular, we can compute the *semantic coherence*: the Pearson correlation between p_i (the influence of node p on node i) and the the negative JSD from node p to node i, J_i. When nodes that are closely related topologically are also closely related semantically (JSD low), coherence is high.

2.4. Community Detection

We expect the links that editors make at the local level to give rise to distinct communities, or *norm bundles*, at the global level. We use the Louvain algorithm [31] to detect global community structure in the Wikipedia norm network. The Louvain algorithm maximizes the modularity at each local partition of the network. The algorithm first assigns each node i to a different community, then computes the potential modularity gain to i for joining the community of its neighbor node j. Each i will join the community of j when the merge offers the highest positive modularity gain. If there is no possible gain in modularity, i remains in its initial community.

3. Results

At first, Wikipedia's population underwent exponential growth. In mid-2007, however, population growth stalled and entered a period of secular decline [26]; see Fig. 1. Over the course of this rapid growth and longer-timescale decay, users created a large number of pages establishing, describing, and interpreting community norms. Our analysis finds a total of 1,976 pages associated with norms. There are 17,235 edges between these nodes; the network density, 0.0044, is of the same order of magnitude as those seen for academic citation networks [32]. There are a total of 56 pages classified as policy, and 113 marked as guideline; for concision, we refer to pages of both types as "policy". The majority of non-policy pages (1,807) are classified as "essays" (1,255), followed by "proposals" (182) (suggestions either rejected by the community, or under discussion) and "humor" pages similar to essays, but taking a more irreverent tone (125).

We were able to achieve good, but not perfect, agreement in categorizing pages as user-content, user-user, or user-admin norms. Our categorization agreement rate was 75% over forty randomly-selected pages. This is well above chance ($p \ll 10^{-3}$); the Cohen's κ value, of 0.59, is on the boundary between "moderate" and "substantial" agreement [33].[1] In the same sample of forty random pages, we encountered only one that we believed was not a norm, giving an approximate precision rate of 97.5%.

[1] We disagreed, for example, on "Editors_should_be_logged-in_users_(failed_proposal)" (user-user vs. user-content) and "Paid_editor's_bill_of_rights" (user-user vs. user-admin).

3.1. Network Construction

Most policy pages appear before the bulk of the population arrives: over half the policy pages are created by May 2005, before the population reaches 20% of its maximum. By the time the population does reach its maximum, in March of 2007, over 80% of the policy pages have already been created. By contrast, the creation of non-policy pages in the form of essays and commentary lags population growth. When the population has reached its March 2007 maximum, less than one-third of the non-policy pages are in place. It is not until a year later that half of the policy pages are in place. This is shown in Fig. 1.

Eigenvector centrality leads to a distinct hierarchy of pages, with some gaining a significant fraction of the overall attention in the system. This is shown in Fig. B.1, broken out by four main page categories—policies, guidelines, essays, and proposals. Policies and guidelines dominate the system by centrality.

The hierarchy is established early, and yet is remarkably stable over the lifetime of the system. The Pearson correlation between the eigenvector centrality of nodes in 2001 and their final values in 2015 is 0.87; year to year, it is always greater than 0.9. The growth in node in-degree is roughly multiplicative; for nodes with degree less than one-hundred (93% of the total network), the growth rate is, on average, $12.7 \pm 0.3\%$ from one year to the next. There is some evidence for super-multiplicative returns to scale; the yearly growth rate for pages with in-degree less than ten is $10.6 \pm 0.4\%$. All of this means that, as new pages enter the system, they fail to gain the prominence of what has come earlier. This leads to an increase in overall network inequality, shown in Fig. 2.

In short, policy growth precedes population growth. Policies have far greater centrality in the network. Attention in the network is unequally distributed and only becomes less equal over time.

3.2. Core Norms

Table 1 lists the top twenty pages in our network. These core norms govern a range of behaviors, including user-content actions (write articles from a neutral point of view, #1; include only verifiable information, #2, and reliable sources, #3), user-user actions (find consensus, #6, assume good faith, #11; be civil, #16; don't "edit war", #19), and user-admin relationships involving specially-defined roles (blocking policy, #13; the arbitration committee, #17). In some cases, a norm spans multiple classes; "What Wikipedia is not", for example, includes both "Wikipedia is not a dictionary" (a norm on the nature of the content to be included) and "Wikipedia is not a battleground" (a norm on how users should interact with each other).

All of these core norms are created early in the system's history. The majority are created before 2004, when the population is less than 3% of the March 2007 peak. The earliest members of the society first define and articulate its core norms.

3.3. Overlap and Semantic Coherence

Over the course of network construction, core norms are drawn apart topologically. At the same time, the semantic coherence of their neighborhoods rises.

Fig. 3 shows the average pairwise overlap between the top ten pages in our network (since some norms are created later, the number of norms in this final set is lower early on). Early in the system history, when the network is small, overlap is very high. The creation of new pages leads to a rapid decline in overlap; even in 2006, when all core norms are in place, the overlap continues to decline. Fig. 3 also shows the evolution of semantic coherence, which rises rapidly and stabilizes early.

Network growth could have been imagined to drive a knitting together of distinct principles. Instead, the opposite happens: core norms slowly draw apart as page creation leads to distinct spheres of influence. Rather than nucleating around a set of densely-connected core principles, the norm network continues to condense around multiple points.

We note that the local clustering coefficient, a measure of the extent to which two nodes, linked to the same node, tend to also link together—remains essentially constant over the span of the data (see Appendix, Fig. D.1). The ways in which editors link together small groups of pages changes little while, their cumulative effect produces to large and lasting changes both in attention inequality and page overlap.

3.4. Emergent Clusters

The network partitions into 14 communities. The four largest comprise more than three-quarters of the network; in Table 2, we describe the top eight, which together compose 97% of all the nodes in the system. By inspecting the top ten nodes in each cluster, we classify these communities into user-content, user-user, and user-admin norms. A force-directed layout (ForceAtlas2, implemented in Gephi [34]), allows us to visualize the norm network and the topological relationships between its emergent communities.

While our network has a fundamentally non-spatial topology at the fine-grained level, the four main clusters arrange themselves roughly in a two-by-two grid. The first dimension moves from interpersonal norms to tasks associated with article writing; top-left to bottom-right. The second dimension moves from interpretive rules (such as how to collaborate well) to technical procedures (such as how to file a complaint); top-right to bottom-left.

Each of the top eight clusters is associated with a distinct topic in our topic model (see Appendix, Table 3); while the article quality cluster is the largest by node number, the topic associated with the collaboration cluster dominates the system by word. Even task-based norms appear to draw on the semantics of interpersonal cooperation.

4. Discussion

The most influential pages in the norm network are also the earliest to be created. We find a Matthew effect [35] for social norms, where later additions to the network do not grow in influence quickly enough to destabilize the hierarchy. We find no normative revolutions on Wikipedia. Why not?

Perhaps the earliest users know best: their policies work well, and are simply adopted by those who come later. Or, later users may join precisely because they subscribe to the norms that have already been articulated. Users who disagree with these norms may find that reinterpretation, rather than replacement, is a more effective response given the disproportionate allocation of attention to early pages.

The fact that core norms are created so early means that a relatively small number of users set them in place. This group may have created norms that meet their own needs, but not the needs of those who arrive later.[2] Recent work [36] has suggested that early users later form an oligarchy that monopolizes power, subverts democratic control, and comes into increasing conflict with the larger collective. If this is true, the enduring centrality of their own interests in the norm network may be a source of power.

Much of Wikipedia's network simply coordinates technical practices such as article naming conventions. The most important norms, however, attempt to rationalize the system around universal concepts such as neutrality, verifiability, consensus, and civility. An important insight comes from a theory of bureaucracy and institutionalized organization developed by Meyer and Rowan (1977 [24]). They propose that norms such as these can function as institutional myths that make the system appear legitimate, and less *ad hoc*, by connecting it to a rational framework.

Page creation continues to grow long after the core norms are already in place. What happens when editors continue to develop and refine this network?

Meyer and Rowan's theory predicts the phenomenon of decoupling, driven by the emergence of inconsistencies between different myths. The essay Civil_POV_pushing, for example, describes how some users may be able to violate the neutrality norm by strict adherence to norms of civility. In Meyer and Rowan's theory, pages like these, that attempt to resolve inconsistencies between myths, will be rare. Myths will instead tend to decouple from each other over time.

Our quantitative findings are consistent with this prediction. As the system grows, the creation of norm-spanning pages such as Civil_POV_pushing are rare and insufficient to prevent the neighborhoods of the core norms drawing apart into separate spheres of influence with high internal semantic coherence.[3] Our findings are also consistent with Meyer and Rowan's second major prediction: that systems become increasingly reliant on a logic of good faith rather than following procedure. Not only is "Assume good faith" itself a core norm, but the associated topic dominates the system as a whole.

The norm network we study here is the culmination of over thirty thousand edits. We analyze the development of this system over time via the editing community's collective decisions and their allocation of attention within the network. While this method tells us a great deal about the collective process of norm creation, we do not know how individual editors understand the relationships between norms or use them to guide how they edit and interact with others. Rather than memorize the complex network in its entirety, an editor may course-grain its properties to form her own mental representation of the encyclopedia's normative structure. Editors' mental representations might then inform their linking and editing behaviors, creating a feedback loop between the representation and the norm network as a whole.

---

[2] If early users are predominantly university students with flexible working hours, for example, they may develop norms that implicitly rely on the possibility of responding to criticism in short, rapid bursts. If later arrivals do not have the same flexibility, but the norms persist, they will find themselves at a relative disadvantage in conflicts that arise—even if the amount of effort they devote to the system each week is the same.

[3] In successful systems, decoupling is also expected to happen not only between myths, but between these myths and actual practice, a phenomenon pointed to by the existence of the page "Ignore_all_rules" ("if a rule prevents you from improving Wikipedia, ignore it").

5. Conclusions

Norms are a crucial unit of cultural evolution, and they gain meaning and force from the relationships that connect them. Our work here has studied the evolution, over fifteen years, of the interdependent network of norms at the center of Wikipedia.

The evolution of this network is a remarkably conservative process. Early features are maintained, and in some cases even amplified, over the course of the network's development. This is consistent with other findings of an "iron law" of oligarchy in peer-production systems.

The encyclopedia's core norms address universal principles such as neutrality, verifiability, civility, and consensus. The ambiguity and interpretability of these abstract concepts may drive them to decouple from each other over time. Wikipedia is a paradigmatic example of a 21st Century knowledge commons. Yet its core norms play a structural role analogous to the institutional myths of rationalized 20th Century bureaucracies.

Acknowledgments

We thank John Miller (Carnegie Mellon), Stephen Benard (Indiana University), and Cris Moore (Santa Fe Institute) for helpful discussions, and the Santa Fe Institute for their hospitality when this work was begun. B.H. was supported by an REU at the Santa Fe Institute under National Science Foundation Award #ACI-1358567, by the Cox Research Scholarship Program, and by the Indiana University Science, Technology, and Research Scholars (STARS) program.

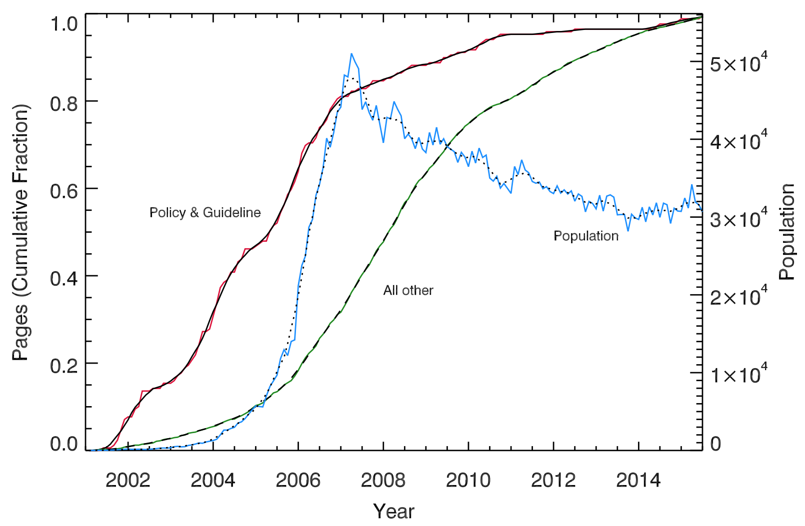

Figure 1. Cumulative growth in policy (red/solid line) and non-policy (green/dashed line) pages, overlaid on active population (blue/dotted line). Policy creation precedes the arrival of the majority of users, while the creation of non-policy pages, usually in the form of essay and commentary, lags the growth in population.

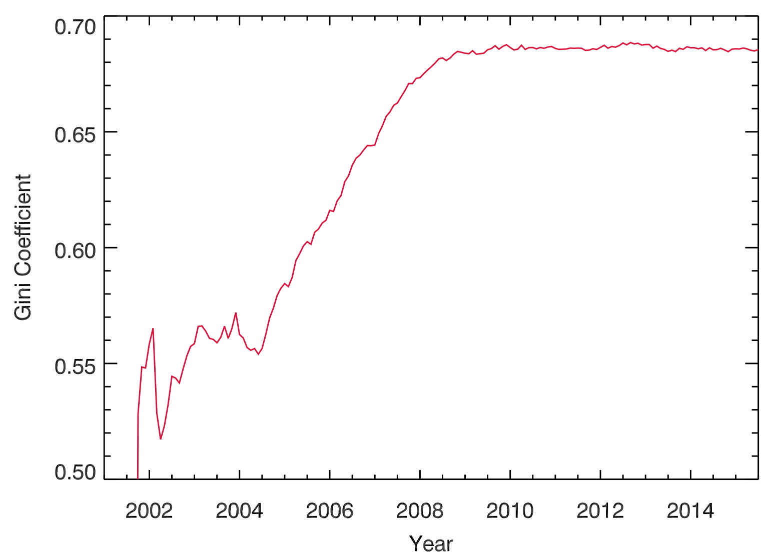

Figure 2. Evolution of the Gini coefficient over time. As new pages enter the system, overall network inequality increases, stabilizing in 2008.

Table 1. Core norms. Top twenty pages, by eigenvector centrality, in 2015. All are either policy or guideline pages, and all are in place by the end of 2006. The majority of these core norms are created before 2004, when the population is less than 3% of its peak.

Rank	Name	Classification	Creation Date
1	Neutral_point_of_view	user-content	24 December 2001
2	Verifiability	user-content	2 August 2003
3	Identifying_reliable_sources	user-content	28 February 2005
4	What_Wikipedia_is_not	user-user/user-content	24 September 2001
5	Biographies_of_living_persons	user-content	17 December 2005
6	Consensus	user-user	11 July 2004
7	Policies_and_guidelines	user-user/user-content	1 November 2001
8	Administrators	user-admin	16 May 2001
9	No_original_research	user-content	21 December 2003
10	Citing_sources	user-content	19 April 2002
11	Assume_good_faith	user-user	3 March 2004
12	Notability	user-content	7 September 2006
13	Blocking_policy	user-admin	8 June 2003
14	Dispute_resolution	user-user/user-admin	12 January 2004
15	Redirect	user-content	25 February 2002
16	Civility	user-user	5 February 2004
17	Arbitration_Committee	user-admin	16 January 2004
18	Vandalism	user-content	29 March 2002
19	Edit_warring	user-user	26 April 2003
20	Talk_page_guidelines	user-user	15 April 2005

Table 2. Top eight Louvain communities, by number of nodes. Communities fall into three classifications (user-user, user-content, user-administration), based on the interactions they govern; we determine these labels by inspecting the top twenty nodes by centrality within each community.

Rank	Fraction of System	Classification	Topic
1	25.9%	User-Content	Article Quality
2	20.9%	User-User	Collaboration
3	15.2%	User-Content	Content Policies
4	15.0%	User-Administration	Administrators
5	8.6%	User-Content	Formatting Articles
6	5.7%	User-User	Wiki-Larping
7	2.9%	User-User/User-Content	Contention & Polarization
8	2.8%	User-User/User-Content	Sensitive Content & Disclaimers

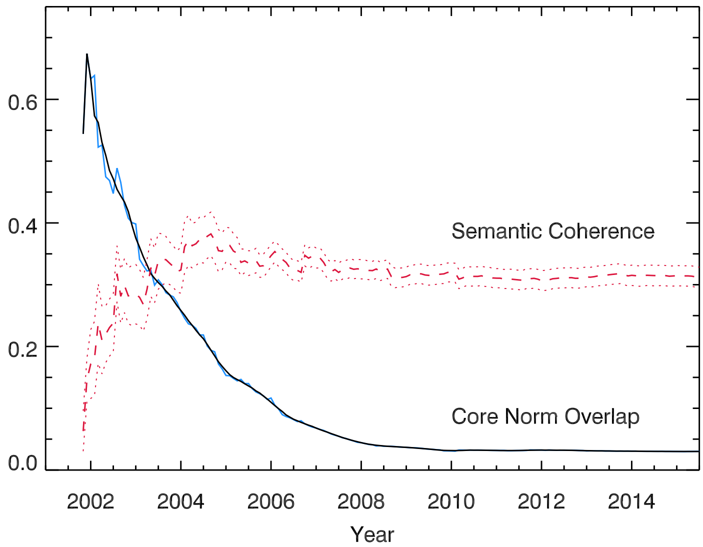

Figure 3. Evolution of influence overlap among the top-twenty norms (by eigenvector centrality) over time (solid line, labeled). In terms of the pages they influence, norms draw apart over time, stabilizing in 2008. At the same time, semantic coherence (dashed line, labeled) increases. Neighborhoods become topologically distinct, but internally coherent.

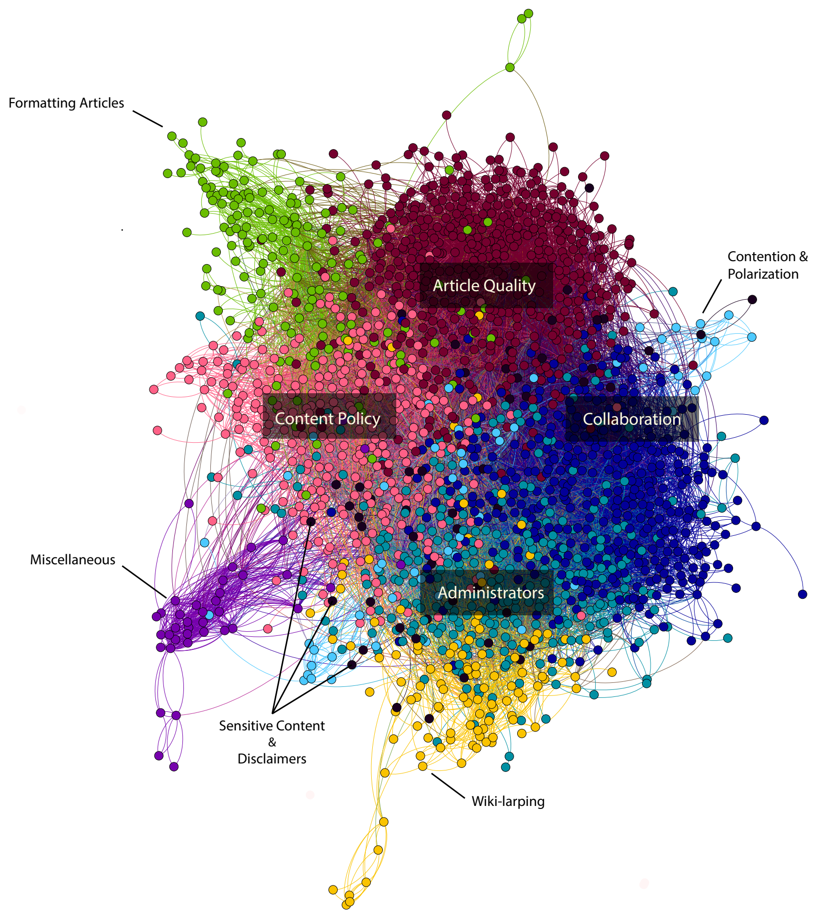

Figure 4. The topology of the norm network is organized into four large, central communities, as found by Louvain clustering. Smaller communities tend to branch from this core. Community themes are based on a sample of high-EC nodes in each community, and confirmed by reference to a topic model based on word-usage.

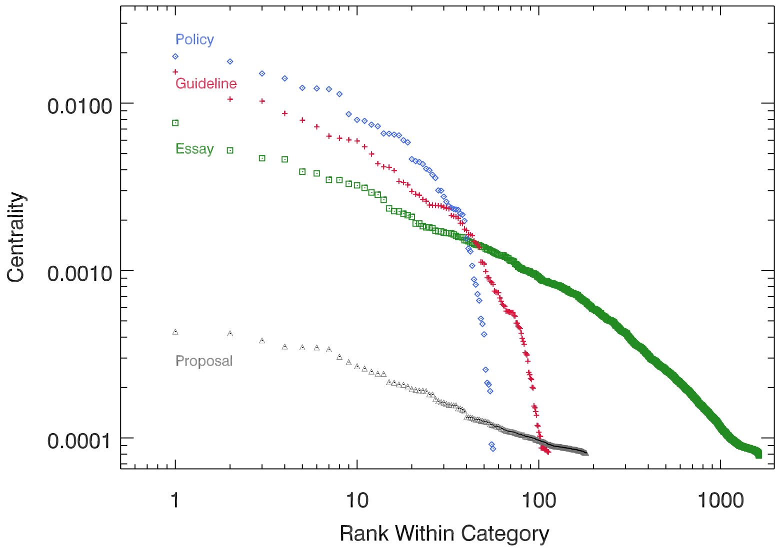

Figure B.1. Ranked eigenvector centrality for pages, broken out by page class. Policy (blue diamond) and guideline (red plus) pages dominate the system. More interpretive essays (green squares; includes humor and related pages), the most common by number, appear at lower relative rank—the highest ranked essay, for example, has lower centrality than the 10th ranked policy. Proposals, failed or current (grey triangles), are the lowest ranked of all.

A. Corpus Construction

As described in the main text, we build our corpus by spidering outward from the page "Assume good faith", following all links in the Wikipedia namespace. Not all pages within the namespace are normative, however. After completing the spidering process, we remove pages that are solely lists (*e.g.*, the pages "List of guidelines" or "Lists of protected pages"), that describe "projects" or other initiatives focused on outreach (*e.g.*, "Wikipedia Loves Libraries"), or on adding a certain kind of content to the encyclopedia (*e.g.*, "WikiProject Libertarianism"), or that serve as noticeboards (*e.g.*, the "Village pump", "Media copyright questions"), with filters on both page titles and editor-assigned categories.

Many page names have synonyms (*e.g.*, "AGF" redirects to "Assume good faith"); we merge synonyms. Not all links between pages indicate a deliberate decision to connect one norm to another. Many pages, for example, contain "boxes", small code snippets that categorize pages or provide navigation indices to similar norms. These boxes can be created by a single command and are replicated across multiple pages; we do not include out-bound links found in these boxes. Finally, pages sometimes have internal links; we drop all self-edges. Our spidering includes only pages that existed on 12:00:00 UTC, 20 August 2015.

B. Combined Scree Plot

Fig. B.1 shows the rank distribution of EC, page by page, broken out by page class. Defining E_i as the eigenvector centrality of the i-th ranked norm allows us to define the break size, $E_i - E_{i+1}$, between this norm an the next. Ranking break-sizes allow us to note positions where the remainder of the norms

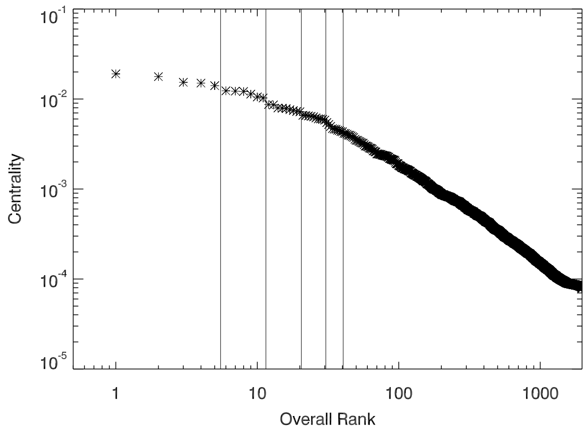

Figure B.2. Eigenvector Centrality for all the pages in our data, ordered by rank. Major divisions (see text) are marked by vertical lines.

in the system have significantly lower EC. Constraining breaks to be greater than five pages apart leads to the top five divisions shown in Fig. B.2. In the main paper, we list nodes up to the third break-point.

C. Gini Coefficient

The Gini Coefficient is a measure between 0 and 1 (inclusive) that calculates the level of inequality in a distribution by comparing it against the *line of perfect equality*, or the point at which 1% of the population holds 1% of the resources, 50% of the population holds 50% of the resources, and so on. To calculate this we use

$$G = \frac{A}{A+B} \tag{2}$$

to find the ratio of the area between the line of perfect equality and the actual distribution of EC, where A is the distance between the Lorenz curve of our data and B is the distance under the Lorenz curve.

Though this measure has traditionally been used to compare income inequality and disparity between countries, its lack of sensitivity to attributes of the distribution such as size make it ideal for measuring inequality even in relatively small distributions.

D. Local Clustering Coefficient

Our work here focuses on the evolution of global network properties, such as eigenvector centrality, overlap, and semantic coherence, that can not be known by breaking the graph into subgraphs. It is interesting to consider more local measures, however, since these are likely to be under far greater direct user control. The example we consider here is average local clustering, defined as

$$\kappa = \sum_{i \in G} \frac{\sum_{j,k \in \mathcal{N}(i)} \delta_{jk}}{|\mathcal{N}(i)||\mathcal{N}(i) - 1|}, \tag{3}$$

or, in words, the number of edges connecting nodes in the neighborhood of i, as a fraction of the total number of possible connections between those neighbors. If individuals have a tendency to connect up the network when they create a new node, by linking together nodes it links to, this will tend to increase

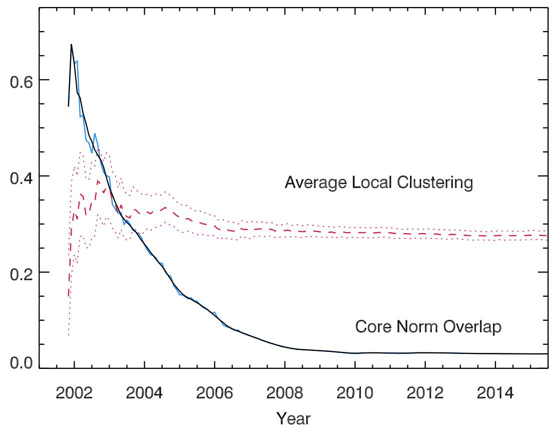

Figure D.1. The average local clustering coefficient, as a function of time. Despite large-scale changes in overall network properties, this local property remains remarkably constant.

the clustering. Fig. D.1 shows this over time. Despite large changes in both population and network size, clustering remains surprisingly constant, at around one-third.

E. Topic Modeling

For our base model with $k = 20$ topics, Table 3 shows the top twenty representative words for each topic; we drop the word "wikipedia", plurals (except the word "wikipedias"), and date/time terms ("january", "utc", etc.). We use Jason Adams' lda-ruby package (https://github.com/ealdent/lda-ruby), a ruby wrapper for the C code of David M. Blei; this code estimates model parameters using variational EM (http://www.cs.princeton.edu/~blei/lda-c/ [29]).

We order the list of topics by their (word-level) prevalence within the encyclopedia. Examining top topics for the eight Louvain clusters described in the main text, we find that each cluster is associated with a distinct dominant topic; these eight topics are noted in column three of Table 3. We note that inspection of the representative words for these eight topics provides complementary evidence in favor of the community labels chosen on the basis of manual inspection of the top ten pages by eigenvector centrality. While the Article Quality community dominates the system by node number, the topic associated with the Collaboration community dominates the system by word.

Rank	Fraction	Louvain Community	Representative words
1	11.4%	Collaboration	editor, edit, dont, good, people, make, editing, policy, page, talk, time, article, faith, point, policies, encyclopedia, consensus, community, personal, user
2	8.67%	Article Quality	source, reliable, article, material, information, research, primary, view, original, editors, subject, published, secondary, policy, neutral, point, scientific, content, topic, claims
3	8.56%	—	article, deletion, page, deleted, discussion, content, delete, speedy, talk, tag, subject, information, policy, user, guidelines, criteria, notability, afd, time, essay
4	8.26%	—	article, information, content, encyclopedia, editors, people, wikipedias, subject, featured, quality, good, list, topic, readers, time, work, project, knowledge, number, lead
5	6.65%	Administrators	consensus, policy, discussion, community, process, committee, arbitration, editors, administrator, user, request, policies, admin, block, dispute, page, wikimedia, proposal, information, made
6	5.80%	Formatting Articles	article, names, title, page, english, disambiguation, naming, redirect, conventions, common, term, style, citation, word, language, topic, book, usage, examples, cases
7	5.69%	—	user, edit, page, vandalism, account, ip, editing, talk, editors, bot, address, protection, administrators, userboxes, username, blocked, block, request, sock, template
8	5.36%	—	notable, article, notability, list, sources, coverage, criteria, information, subject, reliable, emojif, guideline, film, event, university, significant, general, topic, independent, inclusion
9	5.03%	—	page, link, text, image, file, wikimedia, search, commons, web, information, external, software, content, article, site, add, click, wiki, edit, make
10	4.38%	—	talk, edit, page, user, war, im, article, dont, people, time, contribs, good, contributions, back, long, list, things, make, day, ive
11	4.04%	Content Policies	copyright, image, public, nonfree, free, work, content, license, domain, law, fair, article, copyrighted, published, states, pma, united, subject, permission, media
12	3.81%	—	page, talk, template, namespace, user, link, article, text, category, section, special, edit, title, list, signature, ut, mediawiki, redirect, move, navbox
13	3.28%	Contention & Polarization	list, chart, people, united, war, town, world, man, england, states, british, top, hot, songs, women, city, ireland, music, number, death
14	3.04%	—	category, day, categories, article, tip, stub, list, page, people, categorization, main, link, year, created, red, featured, create, template, sort, subcategories
15	2.97%	Wiki-Larping	people, user, time, status, wikidragon, truth, wikifauna, wikipuma, credentials, names, editathon, work, turkish, years, page, make, real, history, group, greek
16	2.83%	Sensitive Content & Disclaimers	support, oppose, policy, people, user, proposal, talk, userboxes, dont, image, offensive, pov, namespace, page, content, article, censorship, vote, npov, agree
17	2.79%	—	ban, topic, editing, indefinite, talk, sanctions, article, page, user, edit, discussion, banned, paid, related, editor, contribs, interest, coi, community, broadly
18	2.55%	—	quotation, style, citing, punctuation, american, mos, ads, dash, manual, inactive, en, english, issue, sentence, dashes, election, text, space, british, jumped
19	2.41%	—	text, template, page, line, article, gt, section, lt, enforcement, table, footnote, law, summary, infobox, style, agencies, synth, color, work, data
20	2.32%	—	article, station, number, year, state, route, highway, time, road, points, date, railway, britannica, ship, include, information, eb, county, class, official

Table 3. Representative words from each of the topics in our $k = 20$ topic model, ranked by the weighted fraction of words assigned. The top eight Louvain clusters are each dominated by a unique topic; these are labelled in column three.

References

1. Sherif, M. *The Psychology of Social Norms*; Harper: New York, NY, 1936.
2. Durkheim, E. *The Rules of Sociological Method*; Free Press: New York, NY, 1938.
3. Akerlof, G. The economics of caste and of the rat race and other woeful tales. *The Quarterly Journal of Economics* **1976**, *90*, 599–617.
4. Geertz, C. Thick description: Toward an interpretive theory of culture. In *Readings in the philosophy of social science*; Martin, M.; McIntyre, L.C., Eds.; MIT Press: Cambridge, MA, 1994; pp. 213–231.
5. Ellickson, R.C.; Ellickson, R.C. *Order without law: How neighbors settle disputes*; Harvard University Press, 2009.
6. Bowles, S. *Microeconomics: Behavior, Institutions, and Evolution*; Princeton University Press, 2009.
7. Simon, H.A. A formal theory of the employment relationship. *Econometrica: Journal of the Econometric Society* **1951**, *19*, 293–305.
8. Tyler, T.R. Psychological perspectives on legitimacy and legitimation. *Annual Review of Psychology* **2006**, *57*, 375–400.
9. Tyler, T.R.; Fagan, J. Legitimacy and cooperation: Why do people help the police fight crime in their communities. *Ohio State Journal of Criminal Law* **2008**, *6*, 231.
10. Elias, N. *The Civilizing Process: Sociogenetic and Psychogenetic Investigations*; Wiley, 2000. Second Edition of 1939 text, edited by Dunning, E., Goudsblom, J., and Mennell, S.
11. Pinker, S. *The Better Angels of Our Nature: Why Violence Has Declined*; Penguin Group, USA, 2011.
12. Klingenstein, S.; Hitchcock, T.; DeDeo, S. The civilizing process in London's Old Bailey. *Proceedings of the National Academy of Sciences* **2014**, *111*, 9419–9424.
13. Ehrlich, P.R.; Levin, S.A. The evolution of norms. *PLoS biology* **2005**, *3*, 943.
14. Ostrom, E.; Hess, C. A framework for analyzing the knowledge commons. In *Understanding Knowledge as a Commons*; Hess, C.; Ostrom, E., Eds.; MIT Press: Cambridge, MA, 2006.
15. Benkler, Y. *The wealth of networks: How social production transforms markets and freedom*; Yale University Press: New Haven, Connecticut, 2006.
16. Bollier, D. The Growth of the Commons Paradigm. In *Understanding Knowledge as a Commons*; Hess, C.; Ostrom, E., Eds.; MIT Press: Cambridge, MA, 2006.
17. Frischmann, B.; Madison, M.; Strandburg, K. *Governing Knowledge Commons*; Oxford University Press, 2014.
18. Ostrom, E. *Governing the commons: The evolution of institutions for collective action*; Cambridge University Press, 1990.
19. Hess, C.; Ostrom, E. *Understanding Knowledge as a Commons: From Theory to Practice*; MIT Press: Cambridge, MA, 2011.
20. Henrich, J.; Boyd, R.; Richerson, P.J. Five misunderstandings about cultural evolution. *Human Nature* **2008**, *19*, 119–137.

21. Shirky, C. *Here comes everybody: The power of organizing without organizations*; Penguin, 2008.
22. Konieczny, P. Governance, Organization, and Democracy on the Internet: The Iron Law and the Evolution of Wikipedia. *Sociological Forum* **2009**, *24*, 162–192.
23. Konieczny, P. Adhocratic governance in the Internet age: A case of Wikipedia. *Journal of Information Technology & Politics* **2010**, *7*, 263–283.
24. Meyer, J.W.; Rowan, B. Institutionalized organizations: Formal structure as myth and ceremony. *American Journal of Sociology* **1977**, *83*, 340–363.
25. Cohen, J.; others. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement* **1960**, *20*, 37–46.
26. Halfaker, A.; Geiger, R.S.; Morgan, J.T.; Riedl, J. The rise and decline of an open collaboration system: How Wikipedia's reaction to popularity is causing its decline. *American Behavioral Scientist* **2012**, *XX*, 1–25.
27. Wikipedia Statistics: Active wikipedians. https://stats.wikimedia.org/EN/TablesWikipediansEditsGt5.htm. Last Accessed: 2015-08-21.
28. Brush, E.R.; Krakauer, D.C.; Flack, J.C. A family of algorithms for computing consensus about node state from network data. *PLoS Computational Biology* **2013**, *9*, e1003109.
29. Blei, D.M.; Ng, A.Y.; Jordan, M.I. Latent dirichlet allocation. *Journal of Machine Learning Research* **2003**, *3*, 993–1022.
30. DeDeo, S.; Hawkins, R.X.; Klingenstein, S.; Hitchcock, T. Bootstrap methods for the empirical study of decision-making and information flows in social systems. *Entropy* **2013**, *15*, 2246–2276.
31. Blondel, V.D.; Guillaume, J.L.; Lambiotte, R.; Lefebvre, E. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* **2008**, *2008*, P10008.
32. Yan, E.; Ding, Y. Scholarly network similarities: How bibliographic coupling networks, citation networks, cocitation networks, topical networks, coauthorship networks, and coword networks relate to each other. *Journal of the American Society for Information Science and Technology* **2012**, *63*, 1313–1326.
33. J. Richard Landis, G.G.K. The Measurement of Observer Agreement for Categorical Data. *Biometrics* **1977**, *33*, 159–174.
34. Jacomy, M.; Venturini, T.; Heymann, S.; Bastian, M. ForceAtlas2, a Continuous Graph Layout Algorithm for Handy Network Visualization Designed for the Gephi Software. *PLoS ONE* **2014**, *9*, e98679.
35. Merton, R.K. The Matthew effect in science. *Science* **1968**, *159*, 56–63.
36. Shaw, A.; Hill, B.M. Laboratories of oligarchy? How the Iron Law extends to peer production. *Journal of Communication* **2014**, *64*, 215–238.