

THIS MANUSCRIPT HAS BEEN ACCEPTED FOR PUBLICATION AT THE JOURNAL OF COMMUNICATION. FOR THE PUBLISHED VERSION AS WELL AS CITATION INFORMATION, PLEASE CHECK: <http://mako.cc/academic/>

Laboratories of Oligarchy?

How The Iron Law Extends to Peer Production

Aaron Shaw (aaronshaw@northwestern.edu)
NORTHWESTERN UNIVERSITY

Benjamin Mako Hill (makohill@uw.edu)
UNIVERSITY OF WASHINGTON

Abstract: Peer production projects like Wikipedia have inspired voluntary associations, collectives, social movements, and scholars to embrace open online collaboration as a model of democratic organization. However, many peer production projects exhibit entrenched leadership and deep inequalities, suggesting that they may not fulfill democratic ideals. Instead, peer production projects may conform to Robert Michels’ “iron law of oligarchy”, which proposes that democratic membership organizations become increasingly oligarchic as they grow. Using exhaustive data of internal processes from a sample of 683 wikis, we construct empirical measures of participation and test for increases in oligarchy associated with growth in wikis’ contributor bases. In contrast to previous studies, we find support for Michels’ iron law and conclude that peer production entails oligarchic organizational forms.

INTRODUCTION

Commons-based peer production – the distributed creation of freely accessible information resources through the mass aggregation of many contributions – represents a model of collective action and public goods production that integrates the use of digital communication networks and information technologies (Benkler, 2006). Peer production has generated public goods of enormous economic, cultural, and political value – such as Wikipedia and GNU/Linux – and has transformed the way that firms in many industries do business, shifted how politicians campaign for office, and changed the way that individuals share in-

formation (e.g., Benkler, 2006; Karpf, 2012; Shirky, 2008). Much of the excitement around peer production stems from the egalitarian principles linked to its organization, to the underlying network communication technologies at its core, and to suggestions that “communal information goods” (see Fulk et al., 1996) will tend to occur in groups with relatively equitable patterns of contribution and less hierarchical organizational forms (Connolly and Thorn, 1990; Rafaeli and LaRose, 1993).

According to various accounts, peer production projects function as novel forms of participatory organization with a broad democratizing potential (Benkler, 2006; Bennett and Segerberg, 2013; Castells, 2012; Fuster Morell, 2012). In addition, peer production communities have inspired a wave of associations, movements, and theorists to embrace networked technologies as tools for creating participatory democratic organizations. However, we also know that many successful peer production projects exhibit strong inequalities of participation and deeply entrenched leadership. Contributions to peer production follow a “power law” distribution where a small group of participants make an enormous proportion of contributions (Healy and Schussman, 2003; Ortega, 2009; Viégas et al., 2007) and many peer production projects, like Linux and Ubuntu, are explicitly undemocratic, going so far as to describe their leaders as “benevolent dictators for life” (Ingo, 2006; Hill et al., 2008). Although unequal participation and entrenched leaders do not necessarily mean that peer production projects are undemocratic, this evidence suggests that they may reproduce a pattern of behavior more consistent with Robert Michels’ “Iron Law of Oligarchy” (1915), which proposes that as voluntary movements and membership organizations become large and complex, a small group of early members consolidate and exercise a monopoly of power within the organization as their interests diverge from the collective’s.

We offer an empirical test of the assumption, central to previous theory, that large peer production projects resist increasing levels of organizational oligarchy. Using exhaustive longitudinal data of internal processes drawn from 683 wikis that have grown large and complex, we adapt Michels’ iron law to the context of peer production communities and construct a series of hypothesis tests. In contrast to previous studies of communal public goods production in online environments (e.g. Rafaeli and LaRose, 1993) as well as ethnographic findings on Wikipedia from Konieczny (2009), we present quantitative evidence in strong support of Michels’ iron law in peer production. As the peer production communities in our sample attract contributors, a small group of leaders, present at the beginning, tend to consolidate power as their interests diverge from those of other participants. These

findings deviate from much previous literature on online organization, peer production, and communal information goods.

Wikis and Wikia Inc.

Our empirical setting is a large sample of peer production communities engaged in the collaborative creation of *wikis*. The term “wiki” refers to software designed to facilitate the collaborative, asynchronous creation and distribution of textual content. It also refers to the communities that use wiki software and to the products created by these groups (Leuf and Cunningham, 2001). Wikipedia is the most famous example of a wiki, but there are hundreds of thousands of other wikis with different goals, topics, and scopes. Through the enormous success of Wikipedia – one of the five most popular websites in the world – wikis are perhaps the most visible and successful modality of peer production.

Along with free/libre open source software (FLOSS), wikis figured prominently into Benkler’s (2002) original conceptualization of commons-based peer production and are widely cited as an archetype of peer production’s new model of economic production, political participation, and information sharing. Despite their novel form, peer production communities like wikis encompass many characteristics that make them comparable to other kinds of voluntary associations. Many of them articulate explicitly democratic organizational ideals: they attempt to maintain member governance, they rely on attracting new members and leaders from within their membership, and they showcase effective solutions to collective action problems.

We analyze a population of peer production wikis hosted by the for-profit firm *Wikia*. Several aspects of Wikia make it an ideal setting in which to compare organizational governance in peer production communities. Wikia, founded in 2004, sought to apply the Wikipedia model of peer production beyond the education-based scope of the Wikimedia Foundation (the organization that supports Wikipedia). Wikia was founded by Jimmy Wales, Wikipedia’s founder, and Angela Beesley, one of the most active and respected contributors to Wikipedia in its early years. Wikia’s policies, structures, and technologies have been heavily influenced by Wikipedia. Although many firms host wikis (e.g., PBWiki, WikiSpaces, and SocialText), Wikia is unique in that it only hosts publicly accessible, volunteer-produced, peer production projects and never restricts viewership.

Wikia relies on peer production to create the content of its websites. As a result, Wikia does not restrict participation in content contribution except to combat spam or vandalism.

Like Wikipedia, anybody can create an account on any Wikia wiki. The vast majority of these wikis allow contributions even without accounts. Like Wikipedia, FLOSS, and other peer production projects, all wiki content is distributed freely to the public.¹ These factors help ensure that Wikia wikis remain open and accessible, allowing communities to adopt more or less participatory and democratic organizational behavior.

In terms of the content and scope, Wikia wikis vary enormously, addressing popular culture and “fan culture” topics as well as things like subcultures, software, food, fashion, and more. Some of the largest wikis host information about massive multiplayer online video games like Halo or World of Warcraft, television shows like *Lost*, movies and novels like *Lord of the Rings*, and information about the academic job market. Taken together, Wikia wikis constitute one of the largest, broadest, and most inclusive populations of peer production communities.

Peer Production as Democratic Organization

A growing body of communication research seeks to evaluate the possibilities for participatory democracy and collective action in an era of digital networks and online organization (e.g. Bennett and Segerberg, 2012; Bimber et al., 2012; Castells, 2012; Polletta, 2013). A chorus of scholars have described networked collectives, including wikis, as deeply democratic (Benkler, 2006; Earl and Kimport, 2011; Fuster Morell, 2010, 2012; Hess and Ostrom, 2011; Karpf, 2012; Kollock, 1999; Tufekci and Wilson, 2012; Weber, 2004). These claims look to peer production projects as a digital vanguard among networked organizations. One popular narrative proposes that technologically-savvy early adopters using the Internet as a tool for encyclopedia writing and software production discovered powerful new mechanisms for avoiding top-down bureaucratic control without sacrificing the quality of their products or processes (e.g., Benkler, 2006; Castells, 2012; Shirky, 2008).

This optimistic view of peer production finds support in several bodies of theoretical work in communication. First, research has linked the democratizing effects of new communication technologies to the production of specific types of “communal” and “connective” public goods like discretionary databases of public information, libraries, census records, and community bulletin boards (Fulk et al., 1996). By the logic of these claims,

¹As is common in peer production, ownership of the copyright on wiki content remains with the contributors but all material is licensed freely to the public as a condition of contribution. All Wikia content is released under the Creative Commons Attribution-ShareAlike license (the same license that Wikipedia uses) and is made publicly available for download.

the production of communal public goods in digital environments occurs more effectively and efficiently in large, egalitarian collectives (Connolly and Thorn, 1990; Fulk et al., 1996; Rafaeli and LaRose, 1993).² Second, although focused on a social movement context, Bennett and Segerberg’s “logic of connective action” suggests that networked modes of collaborative, voluntary association driven by personalized use of technology-mediated communication – including those pursued in peer production communities – will entail flatter, less hierarchical organizational forms (Bennett and Segerberg, 2013). Finally, within the “collective action space” of Bimber et al. (2012), the members of peer production communities embrace “entrepreneurial” (i.e., less bureaucratic and routinized) modes of engagement along with more “personal” (i.e. less formal and instrumental) modes of interaction. Bimber and colleagues contend that such entrepreneurial-personal organizations tend to be relatively flat and are managed through informal, discursive norms. Across the organizational communication literature, these theories predict that peer production communities ought to possess participatory, democratic characteristics.

Peer Production and Michels’ Iron Law of Oligarchy

At the same time, empirical research on leadership, governance, and participation in peer production communities suggests a much more complicated reality than narratives emphasizing non-hierarchical and non-bureaucratic organization. First, nearly every population of peer production projects studied follows a “power law” distribution of contributions across contributors (e.g. Schweik and English, 2012; Ortega, 2009; Wu et al., 2009). Second, there is some evidence that the informal hierarchies of attention, status, and influence that arise in large, successful, peer production projects cohere into formal and increasingly rigid bureaucratic structures that may not facilitate participatory democracy. For example, radical inequalities of status and participation are pervasive in Wikipedia (e.g. Viégas et al., 2007; Kittur et al., 2007; Loubser, 2010; Priedhorsky et al., 2007; Panciera et al., 2009) and studies of the internal governance of online collectives indicate that hierarchical institutions exist, even if those institutions are not accompanied by formal bureaucratic structures (e.g. Butler et al., 2008; O’Mahony and Ferraro, 2007; Shaw, 2012; Forte et al., 2009; Zhu et al., 2011; Black et al., 2011). While extreme hierarchy, bureaucracy, and inequality do not rule

²More recent studies have turned to laboratory settings as well as firms and non-profit organizations, and have bypassed the question of whether specific modes of organizational governance more effectively support communal information resource production (e.g., Child and Shumate, 2007; Cho and Lee, 2008; Raban and Rafaeli, 2007; Yuan et al., 2007).

out the possibility that peer production projects are democratic, these attributes are not typical of participatory organizations.

A related set of predictions derive from the literature on voluntary organizations as well as some of the previous studies on peer production. The “iron law of oligarchy,” one of the most influential theories of governance in voluntary organizations that was developed by the German sociologist and political theorist Robert Michels (1915), encapsulates these perspectives. Michels sought to explain why democratically-run political parties in Italy around the turn of the twentieth century became less democratic as they grew larger and more complex.³ His “iron law” proposes that two patterns drive increasing oligarchy within voluntary organizations as they grow: (1) structurally, the group becomes an increasingly formal and complex organization with a small group of professional leaders who exercise a monopoly over the mechanisms of authority; and that (2) these leaders develop independent interests in the preservation of the organization itself, resulting in the transformation of the goals and activities of the organization in ways that diverge from the interests of members. Following previous work, we refer to these distinct dynamics as the “structural” and “goal transformation” components of the iron law (Jenkins, 1977; Leach, 2005; Voss and Sherman, 2000).⁴

Previous studies of movement and voluntary organizations – in particular Lipset et al. (1956), Piven and Cloward (1977), Rucht (1999), and Voss and Sherman (2000) – guide our operationalization of Michels’ two components of oligarchy in the context of peer production projects. Structurally, the governance and leadership of oligarchic organizations must reside in the hands of a stable, entrenched, minority that exercises dominant control over organizational resources and policy (Lipset et al., 1956). Oligarchic goal transformation occurs when organizational leaders and elites develop interests that diverge from those of other members and impose their interests on the rest of the organization. If it applies, Michels’ iron law requires that both patterns hold across peer production projects as they grow over time.

Although distinct from oligarchy, broad evidence of hierarchy, bureaucracy, and inequality in peer production suggests oligarchic leadership may prevail. Additionally, un-

³We do not attempt, in this article, to give a comprehensive overview of research into the iron law and its application. In addition to Michels’ original text, we refer readers to work by Jenkins (1977), Leach (2005), and Voss and Sherman (2000).

⁴It is worth noting that the conceptual elaboration of the iron law that we present here follows Jenkins (1977) and Voss and Sherman (2000) more closely than it does Leach (2005).

like political parties or labor unions, wikis, in general, are not formed with the objective of facilitating political or economic representation of their members and, as a result, may be less committed to democratic governance. Finally, given the ease with which volunteer contributors may choose to “exit” the community, some previous work on voluntary associations would also predict that less democratic governance arrangements would prevail (Barakso and Schaffner, 2008).

On the other hand, the organizational form and governance of peer production projects in general, and wikis in particular, possess few technical or physical constraints. Not only does wiki software support widely distributed contributions, it also supports widely distributed authority and governance. For example, there is no limit on the number of formal leaders (“administrators”) a wiki can have and no material cost to adding them. Indeed, in previous research, wikis have been considered some of the most egalitarian and democratic peer production projects in terms of their opportunities for lateral authority and resistance to domination by a small minority (Konieczny, 2009). If peer production projects are, as some suggest, intrinsically supportive of participatory democratic organizational forms, there is reason to believe that wikis will be among peer production’s most democratic organizations.

Testing for the Iron Law in Wikis

Despite the broad interest in peer production as democratic organization, little previous empirical research on peer production has formally tested such claims about projects’ organizational governance and several studies point to the need for more comparative analysis (Butler, 2001; Fuster Morell, 2010; O’Mahony and Ferraro, 2007; Shaw, 2012). Also, even though theoretical and empirical work asserts that organizational democracy and egalitarian participation will prevail in communities engaged in communal public goods production, concepts of organizational democracy are often left unspecified. For our purposes, organizational democracy consists in the active participation of multiple constituencies in the negotiation and exercise of legitimate authority within an organization. This definition derives from previous research on voluntary, membership, and movement organizations engaged in collective action of various kinds (e.g., Andrews et al., 2010; Lipset et al., 1956; Ostrom, 1990; Voss and Sherman, 2000) and builds on foundational studies of public goods production and common pool resource management (e.g., Hardin, 1968; Olson, 1965; Ostrom, 1990). As with other translations of the iron law to organizations beyond political

parties, oligarchy and democracy have distinct meanings in the context of peer production communities. Like previous work, we adapt and operationalize it to suit our empirical setting.

The structural component of the iron law concerns the distribution of formal authority within an organization. Wiki software facilitates role differentiation and systems of formal leadership. In most wikis, members can hold special privileges that include the ability to delete accounts and content, to lock and unlock pages for editing, and to block users. Other special privileges include access to special pages and to tools for removing spam and vandalism. Users with additional privileges are referred to as “administrators” within wiki communities. As individuals occupying structural positions of authority, administrator status provides an avenue to test for the structural component of the iron law. Our first hypothesis tests for structural oligarchy by considering the relationship between the distribution of administrator status and membership size: *The probability of adding new administrators declines as wikis’ contributor bases grow (H1).*

Although administrator status confers authority and special privileges, the practice of governance and leadership in wikis consists largely of the exercise of informal authority and the performance of tasks that loosely resemble the work of research editors and debate moderators (Forte et al., 2009; Zhu et al., 2011; Black et al., 2011). So although only an administrator can delete a page or ban a user, non-administrators can, and frequently do, contribute to administrative pages in ways that may traditionally be considered part of bureaucratic administration of the community (Butler et al., 2008; Joyce et al., 2012). Much of this governance activity happens in dedicated administrative pages that, in most cases, any wiki user can create or edit (Kriplean et al., 2007). Although not every wiki uses these administrative pages, they act as dedicated spaces for crafting rules, policies, and guidelines (e.g., rules for deleting articles) as well as an area in which policies are implemented (i.e., discussion and voting on the deletion of particular pages). Our second hypothesis looks to these pages to consider the relationship between the administrative elite and participation in administrative activity: *Controlling for the total number of contributions to administrative pages, administrators will contribute more to administrative pages as wikis’ contributor bases grow (H2).* This hypothesis also tests for the structural component of oligarchy, and in particular the idea that oligarchic leaders monopolize the mechanisms of organizational authority.

Of course, the presence of a consolidated elite may not be problematic when the elite

represents the interest of the broader membership. Structural indicators of oligarchy, in this sense, only become a problem in the presence of goal transformation when leaders use their positions of authority to alter the priorities of the organization to suit their own agendas over those of the organization’s members. Support for H2, may reflect a shift of activity of the elite into governance and away from the interests of the membership as expressed by how they contribute to the wiki. For more direct evidence of goal transformation, we can examine a subset of “reverts” (i.e., complete removals of contributions) performed by administrators to other editors. Reverts, a technique used to reject low quality, spam, or otherwise unwanted contributions, reflect a way that administrators impose their vision of legitimate and valuable participation on others.⁵ More importantly, reverts provide a clear indication of divergence between the interests of the reverters and the reverted. If a contributor makes a good faith edit and that edit is undone, this signals a difference between the reverting and reverted users in terms of what they believe the content on the website should be. Formally, we predict that: *Controlling for the total number of contributions made by experienced contributors, the number of reverts by administrators of such contributions will increase as wikis’ contributor bases grow (H3).*

METHOD

Data and Measures

In contrast with previous studies assessing either the iron law or the organizational dynamics of communal public goods production, our analysis draws on the full observational records of intra-organizational behavior within a population of organizations. This type of “big data” affords us an opportunity to test hypotheses in ways that previous studies on these topics could not. Our initial dataset includes data on 76,473 wikis created before April, 2010. It includes the full history of every revision, from both registered and non-registered users, for every wiki hosted by Wikia from the time of Wikia’s founding in 2004 until the point of data collection in April, 2010. Because Wikia does not host private wikis or private data, these datasets were systematically made public for every Wikia wiki. Data for all wikis in our sample remains public and searchable on the web and new datasets continue to be available, upon request, for every wiki. Wikia wikis all run the same software

⁵Research on Wikipedia has shown that reverts by administrators tends to drive away newcomers (Hal-faker et al., 2011). Assuming a similar dynamic exists in wikis in our dataset, the presence of reverts by administrators might provide an indicator of the extent to which administrators prevent less experienced contributors from acquiring authority or control of the organization.

Variable	Minimum	Median	Mean	Maximum	SD
Edits	644	13438	53306	2303248	161652
Pages	183	3167	11152	1270640	53025
Editors	69	218	787	68222	3457
Reverted Edits	0	285	1441	122950	5886
Administrators	0	7	11	247	18
Age (Months)	6	46	50	74	11
Project Edits	0	55	622	59726	3224
Experienced User Edits	418	12270	49020	2020925	149606

Table 1: Summary statistics for all of the wikis included in our analysis. ($n = 683$)

(MediaWiki) which is developed in collaboration with the Wikimedia Foundation. The fact that all of the wikis in the population use identical software makes it possible to derive a set of comparative metrics for analyzing organizational governance and activity.

Although scholars often treat all wikis as examples of peer production, there is a highly skewed distribution in activity among wikis; the vast majority of wikis, like FLOSS projects (Healy and Schussman, 2003), are small and uncollaborative. Considering sample median values, an average wiki in our full dataset contained 225 contributions to 217 pages by 6 unique contributors at the point of data collection in 2010. Most of these wikis are not produced through the type of mass collaboration that Benkler (2002, 2006) calls peer production. For this reason, and because both theories and measures of oligarchy assume large, active groups, we limit our analysis to a subset of large wiki communities from this population.

To build this subset, we first ranked all 76,473 wikis in terms of the number of unique contributors. Next, we selected the top one percent of wikis in the full dataset (732 wikis). We removed a number of wikis because their database dumps included corrupted or invalid data (e.g., edits marked as occurring before the wikis were founded) leaving a total of 683 wikis in our final dataset. Summary statistics for this subset of wikis are reported in Table 1. This sample is an inclusive sample of peer production Wikia wikis in that the smallest of these wikis includes contributions from only 69 accounts. Because a one percent cutoff is still arbitrary, we also carried out our analysis using less inclusive subsets of only the top 300 and top 100 wikis. In both cases, we found results that were substantively similar.

This dataset consists of rich longitudinal records which include every contribution made to each wiki, recorded with timestamps accurate to within one second. It includes

33,278,993 distinct contributions by 469,524 different contributors to 6,167,797 different wiki pages – more than 264 gigabytes of raw data. To create the measures in our analysis, we wrote custom software to parse the database dump files and to query the Wikia Application Programming Interface (API) for additional metadata. Because statistical power is not a concern and because discrete longitudinal analysis is more easily interpreted, we “bin” our data into week-long periods for this analysis. As a result, our analytic unit is the *wiki week* and our longitudinal dataset includes 146,858 such observations.

Since Michels’ iron law is stated in terms of membership growth, or organizational complexity, our primary independent variable must capture the size of a project’s membership. For all our hypotheses, our primary question predictor, *accounts_{total}*, measures the number of unique registered accounts that have made at least one contribution to a wiki. Because wikis can grow older as well as more complex – and because our detailed data can disentangle these effects – we include another variable, *week*, that is expressed as the age of each wiki in weeks from the time when the wiki recorded its first contribution. As introduced above, we use “administrator status” as an indication of structural and formal authority of individuals within wikis. We calculate each individual’s administrator status on a weekly basis for every *wiki week* in the dataset. Because this data is longitudinal as well, activity of a contributor in one week may be classified as coming from an administrator, while activity from the same user in another week may not.

To test our three hypotheses, we construct three dependent variables. For our first hypothesis (H1) that the probability of a community adding a new administrator will decline as wikis’ contributor bases grow, we create a dichotomous variable that is coded “1” in a week if a community adds a new administrator and “0” if it does not. The addition of a new administrator constitutes a rare event and only occurred in 2,911 wiki weeks (2% of the wiki weeks in our sample).

To test H2 that structural oligarchy will also be associated with increased administrator activity in governance, we first identify a set of edits to administrative pages in each wiki. We take advantage of a feature of MediaWiki software called “namespaces” which allows pages to be categorized into one of several groups. One namespace common to all wikis is the “project” namespace which, according to documentation, is designed to be used for “meta-discussions related to the operation and development of the wiki”. By convention, this namespace is used for administrative and policy activity including policy and guideline creation and decision-making (Kriplean et al., 2007). As a first step, we identify a subset of

edits to pages in the project namespace and record a count of these edits ($proj\text{-}edits_{week}$) for use as a control. As a second step, we identify the subset of edits made by administrators ($proj\text{-}edits\text{-}admin_{week}$) for use as an outcome. Some wikis in our sample do not use the project namespace (129 wikis or 19% of our sample) so we exclude these wikis from our analyses when testing H2. Our results are substantively unchanged if they are included.

To test H3, that the number of reverts by administrators of edits made by experienced editors will increase as wikis' contributor bases grow, we must first identify reverted edits. We identify reverts as revisions that return an article to its state prior to the last editor's contribution.⁶ As noted above, reverts that administrators make to the contributions of other community members provide a rough measure of the extent to which administrator interests diverge from those members' interests. Many reverts reflect work undoing vandalism and spam. Although it is not possible to automatically categorize contributions as having been made in either good or bad faith, we can consider a contributor's total edit history to identify experienced users with a history of accepted contributions. Toward this end, we construct a variable $admin\text{-}reverts_{week}$ that is a count of the number of edits of experienced editors that were reverted by administrators in each week on each wiki. We limit this count to reverts of edits made by experienced editors because we do not want the measures to reflect reverts of spam and vandalism. We label contributors "experienced" who have made at least five non-reverted edits and have not had more than 25% of their total edits reverted. We also construct a measure of the number of edits by these experienced editors in each week ($expr\text{-}edits_{week}$) for use as a control.

Finally, we construct a baseline set of control variables. To capture differences in activity that may affect the current amount of administrative work, we construct measures of the number of unique contributors ($editors_{week}$) and the number of distinct contributions ($edits_{week}$). We also include controls that capture the state of the wiki including the number of distinct pages in the wiki ($pages_{total}$) and the total number of administrators ($admins_{total}$).

Analytic Strategy

To test the three hypotheses described above, we construct longitudinal models to estimate the extent to which the wikis in our sample become more oligarchic as their contributor

⁶Specifically, we treat a contribution as reverted if the previous revision is identical to the subsequent revision. In other words, we treat a contribution as reverted if, and only if, a user makes a contribution and the next contribution simply undoes her work. Because a contributor's edits can be undone in ways that do not return the page to its previous state, this reflects a conservative measure of reversion.

bases grow. Because our dataset includes repeated observations for each wiki, we are concerned about autocorrelation of residuals over time. Following Singer and Willett (2003), we use hierarchical linear models as a multilevel model for change and fit random intercepts for each wiki to cluster within-wiki variance in a compound error covariance structure.

As is typical in datasets from online communities, many of the variables capturing activities in wikis are highly skewed. Our measures of accounts, editors, pages, administrators, edits, administrator-reverts, edits by experienced editors, and administrative edits all follow this pattern. We take the natural logarithm of each, resulting in approximately normal distributions. Preliminary exploration of our data revealed that the relationship between time and our dependent variables is curvilinear. As a result, we include a quadratic specification of *week* in our models.

Each of our models is fit with a measure of oligarchy as its dependent variable and each model corresponds to one of our hypotheses. In Model 1 (M1), we use a multilevel logistic regression to estimate the probability of a new administrator being added. Models 2 and 3 are hierarchical linear models on different dependent variables: (M2) is the log-transformed number of edits to administrative pages by administrators, and (M3) is the log-transformed number of reverts of experienced users by administrators. We use a base model in which every variable is measured at the level of the wiki week and which includes a set of controls as well as our compound error term:

$$Y = \beta \ln accounts + \beta \ln week + \beta \ln week^2 + \beta \ln pages + \beta \ln admins + \beta \ln editors \beta \ln edits + [u + \epsilon]$$

M2 estimates log edits by administrators on administrative pages and adds a control for the total amount of such activity (*ln proj-edits*). M3 estimates the log number of administrator reverts of experienced contributors' edits and includes a control for the total number of these edits (*ln expr-edits*).

RESULTS

We find that as the wikis in our sample attract contributors, they display characteristics consistent with increases in oligarchy. Leadership roles remain concentrated among an extremely small group of elites as organizational membership grows, affirming H1. This group of administrators is increasingly active in administrative activity, controlling for the




Figure 1: Cumulative plots of covariates for *Seattle Wiki*, a collaborative website for information about Seattle and one of the online communities in our dataset.

total amount of such work, consistent with H2. Finally, we observe that as membership grows, the number of reverts of experienced users by administrators tends to rise, supporting H3. In all three measures, wikis in our sample conform to the patterns predicted by Michels’ iron law.

Figure 1 provides an example of a single wiki from our dataset that reflects the basic pattern demonstrated in our models. The data shown are from *Seattle Wiki*, a website that contains information about, “places to go, things to do, community, organizations, happenings and events, resources, politics, and anything else related to the greater Seattle area.”⁷ Figure 1 plots cumulative counts of several variables. The visualizations show that

⁷<http://seattle.wikia.com/>.

the number of contributions (*edits*), the number of registered users that have made at least one contribution (*accounts*), and the number of edits to “project” pages have each increased steadily from the time that the community was created in late 2004 through the point of data collection in 2010. That said, it is clear in the third panel that the number of administrators has grown much more slowly and that no administrators have been added since late-2005. The bottom panel shows that this small group of administrators has, over time, continued to revert contributions from experienced users, but that this activity has slowed.

Fitted regression models suggest a substantively similar pattern of associations and provide a set of formal tests of our hypotheses. Results from these models are shown in Table 2. In our first model, (M1), we find strong support for H1, that large communities are less likely to add new administrators as they add new members. Indeed, we find that, *ceteris paribus*, one log-unit increase in the number of registered contributors is associated with odds of adding a new administrator which are only 0.81 times as high as they would be with the smaller contributor pool. This estimate reflects a marginal negative effect controlling for the number of administrators, the amount of activity, and the number of individuals active in each week. As expected, a large administrator corps is associated with a lower probability of adding a new administrator while a large number of active editors in a given week are a positive predictor of a larger leadership body.

Our second model (M2) suggests support for H2 that as communities’ contributor bases grow, the number of edits made by administrators to administrative “project” pages will increase. Because both our dependent and independent variables are expressed in natural log transformed units, our parameters can be interpreted as elasticities. At the margin, we estimate that a 1% change in the total number of accounts on a wiki, controlling for the total edits to administrative pages, is associated with a 3% increase in the number of contributions made by administrators to the administrative parts of the wiki.

Our final model (M3) suggests support for H3 that as communities’ contributor bases grow, the number of edits made by experienced contributors that are reverted by administrators also tends to grow. Again, the effect can be interpreted as an elasticity. At the margin, we estimate that a 1% change in the total number of accounts on a wiki, controlling for the total number of edits by experienced users in the wiki that week, is associated with a 5% increase in the number of contributions made by experienced editors being reverted by administrators.

	M1	M2	M3
(Intercept)	-4.108*** (0.081)	-0.138*** (0.011)	-0.029* (0.015)
<i>week</i>	-0.006*** (0.001)	0.000*** (0.000)	-0.002*** (0.000)
<i>week</i> ²	0.000*** (0.000)	0.000*** (0.000)	0.000*** (0.000)
ln <i>accounts</i> _{total}	-0.210*** (0.031)	0.025*** (0.002)	0.045*** (0.002)
ln <i>pages</i> _{total}	-0.763*** (0.023)	-0.009*** (0.001)	-0.021*** (0.002)
ln <i>admins</i> _{total}	0.666*** (0.035)	0.070*** (0.005)	-0.010 (0.006)
ln <i>editors</i> _{week}	0.116** (0.038)	-0.036*** (0.002)	0.125*** (0.003)
ln <i>edits</i> _{week}	0.996*** (0.022)	0.010*** (0.001)	-0.111*** (0.005)
ln <i>proj-edits</i> _{week}		0.608*** (0.002)	
ln <i>expr-edits</i> _{week}			0.178*** (0.004)
Log Likelihood	-11034.750	-38022.642	-103351.946
Num. obs.	146858	118994	146858
Num. groups: wiki	683	554	683
Variance: 1 wiki	0.325	0.034	0.080
Variance: Residual	1.000	0.109	0.234

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table 2: Table of fitted multilevel regression models. The unit of analysis in each case is the wiki week. *M1* is a logistic regression regression model of the probability that a wiki will add a new administrator during a week. *M2* is a linear model predicting the logged number of edits made by administrators on administrative “project” pages controlling for total edits to these pages. *M3* is a linear model predicting the logged number of reverts of edits by experienced editors by administrators controlling for the number of edits by experienced editors. ($n = 146,858$ wiki weeks from 683 wikis.)

To assist in further interpretation and comparison, plots of predicted values for prototypical wikis for each of our three models are shown in Figure 2 along with 95% confidence intervals for the marginal effects. These model-derived plots emphasize what happens, on average, in our sample. The values along the y -axis correspond to estimated values of each of our dependent variables in our three models. Along the x -axis of each plot are a range of




Figure 2: Plots showing predicted values from our models for wikis with varying number of accounts holding all other variables at sample medians. The graph also includes 95% confidence intervals for the marginal effects using the methods and tools described in Fox (2003). All outcome variables are measured in “per week” units.

likely values of total registered accounts from 0 to 793 (the 95th percentile of observations in our dataset). These plots show the predicted values for wikis that differ in terms of the number of registered users with at least one edit but are identical in every other respect. We have held each of our control variables constant at the sample median. As is the case in our models, each of these prototypical values should be interpreted in the context of a single week-long period. One important limitation is that each of these plots controls for time. As organizations grow larger and more complex, they also grow older. Although corresponding plots over time are not shown here, we find that the effect of age will, in most cases, compound the effects of organizational size. For example, older wikis are usually less likely to add administrators even once we hold the membership size constant.

Establishing the substantive significance of these results presents challenges for several reasons. Although M1 suggests new administrators become much less common in wikis with more contributors, the size of the estimated marginal effects in M2 and M3 for the

prototypical wiki in Figure 2 is at the level of fractions of edits per week, and is estimated to be effectively flat over much of the variation in community size in our sample. Small movements in our indicators of oligarchy may have large effects. For example, Wikipedia researchers have shown that even a single revert can have powerful demotivating effects (Halfaker et al., 2011) with broad implications (Halfaker et al., 2013). Then again, an increase of a single reverted edit every several weeks, or of a slightly larger share of administrative editing falling on administrators, might not be felt as a stark shift into oligarchy. Because of these challenges, and because Michel's iron law makes no attempt to quantify oligarchy, our hypotheses are framed in terms of directionality. Our results reflect consistent movement toward oligarchy in terms of our measures.

DISCUSSION

The pattern of results in Table 2 and Figure 2 provides strong evidence that, on average, as wikis' contributor bases grow, a small group – present at the beginning – monopolizes positions of formal authority in the community and accounts for more administrative activity while also using their authority to restrict contributions from experienced community members. These findings contradict much theoretical and empirical scholarship suggesting that peer production communities will resist oligarchy and embrace participatory organizational practices. In the large sample of peer production communities we examined, wikis tend to reproduce undemocratic, non-inclusive, organizational hierarchies. The wikis in our sample are not indicative of robustly democratic, participatory institutions. This is true despite the relative lack of formal bureaucratic structure or clearly-defined roles within many wikis. These results are consistent with Michels' iron law of oligarchy and contradict prevailing notions regarding organizational democracy in peer production. Of course, our work is also subject to several threats and limitations.

Limitations

One important limitation concerns the effectiveness of the controls for the number of administrative edits (*proj-edits*) in M2 and contributions by experienced editors (*expr-edits*) in M3. Without these controls, more administrative edits or reverts by administrators might simply reflect an increase in administrative work and editing associated with a more active community. Although our controls for baseline activity in these two variables are among the strongest estimated effects in our models, they may still be insufficient. We gain

confidence from several alternative specifications used as robustness checks included as an appendix in Table 3. In Model 4, we alter M2 to estimate the logged count of the number of edits to project pages by *non*-administrators and find no statistically significant relationship to membership size. Controlling for baseline project page activity, administrators increase their editing as wikis grow, while non-administrators do not. An alternative approach to controls is to express our dependent variables as proportions. Models 5 and 6 in Table 3 are specifications of M2 and M3, respectively, using proportions as a dependent variable instead of our activity-level controls. Because proportions require a non-zero denominator, these datasets omit wiki weeks with no administrative edits or edits by experienced editors. These models show substantively similar results to M2 and M3: membership size is positively associated with the proportion of contributions to administrative pages made by administrators and with the proportion of edits by experienced editors that are reverted by administrators.

Another potential threat to the internal validity of our findings concerns a key measure of oligarchic structure. In all three of our tests, our analysis relies on the administrator status as a way of identifying elites. However, as we explain in our background section, the presence or absence of administrator rights are only one way of measuring elites in wikis. As wikis grow, contributors may take on the activity of much wikis governance (i.e., creating, invoking, and enforcing rules) while merely relying on the small pool of administrators to carry out the final act of banning or deleting content. An alternate specification of M2 that predicts the logged number of editors contributing to project pages – ignoring whether they are administrators – can begin to address this concern. Estimates from this fitted model are shown in Model 7 in Table 3. In support of our finding in H2, we estimate a negative relationship between the number of editors contributing to project pages and membership size, controlling for activity in administrative pages. *Ceteris paribus* and administrative rights aside, fewer people are involved in governance activities as wikis attract contributors. Future work might use network measures or heuristics to identify a group of elite users without relying on administrator status.

Our analysis is also limited in ways that are more difficult to address. We employ “big data” methods from software engineering to construct our dataset, we test our hypotheses using more familiar regression-based tools. Although these methods are common in the social sciences, they import parametric assumptions about log-transformed data that can lead from subtle model misspecification to misleading findings. Our choice of methods

was driven by the nature of our research questions, and we have built confidence in our results through alternative specifications and robustness checks, but we cannot eliminate these threats.

The breadth of our analysis and the size of our dataset means our measures do not include some of the nuance possible through other methodologies. Although we attempt to address qualities of individual communities through wiki-level random effects, our analysis is one of averages and cannot speak to differences that exist across the genres, languages, and demographic characteristics of the communities in our sample. By using data intensive techniques to look across millions of contributions and hundreds of communities, we lose the ability to speak in depth about some differences between and within those groups. We have spent time reading and contributing to some of the wikis in our dataset to build an intuition and familiarity with the users and communities in our data. That said, our measures necessarily provide an incomplete picture of a broad and rich dataset. The question of the extent to which the patterns we observe constitute meaningful movement toward oligarchy requires further in-depth analysis of individual projects.

Generalizability

As is always the case with studies of organizations, movements, and communities, important questions about generalizability remain as well. Wikis are a particular type of technological platform and Wikia wikis tend to be used for particular types of projects. As a result, we cannot be certain about the implications of our findings for the broader universe of peer production. There are also important questions about generalizability of our results beyond peer production. Many of the theories of communal information goods and collective action were formulated without peer production or the Internet in mind. Peer production projects are also unusual kinds of organizations that share characteristics with firms, cooperatives, political parties, unions, social movements, fan clubs, self-help groups, and more. As a result, a comparative analysis of peer production employing organizational theory is, ideally, an exercise in careful analogy. At the same time, peer production projects are organizations in the fundamental sense that they are collectivities with goals and boundaries. Moreover, as voluntary organizations, peer production projects like Wikipedia and GNU/Linux have had unprecedented success at mobilizing collective action and contributions to communal public goods. For this reason, we join communication and social movement scholars, as well as proponents of peer production, in the belief that theories of

public goods production and voluntary organizations are a salient reference point for research in organizational democracy. The question of whether or not our findings regarding peer produced communal information goods apply to other sorts of networked movement organizations remains for future research.

For scholars of collective action who are accustomed to environments where the objectives may include things like information shared within or between firms, political rights, fair wages, or representation from the state, it may be tempting to dismiss the forms of leadership and activity on wikis as trivial or insubstantial. We believe that such dismissal is shortsighted given the contemporary influence of peer production technologies and organizational strategies. For example, in Mexico, a newly created political party constructed around a call for participatory democratic political structures calls itself the *Wikipartido*.⁸ Even the most playful peer production communities face many of the same obstacles to collective action and public goods production as other kinds of volunteer-based membership organizations. As such, peer production projects provide an ideal environment within which to test and elaborate existing theories and laboratory-based findings using observational data.

Conclusion

This paper offers several contributions to existing research on peer production and collective action in participatory organizations. First, it joins a small number of studies in expanding the existing domain of organizational research beyond political parties, firms, non-profits, and social movements in order to test influential organizational theories in the context of peer production. Second, it contributes to our understanding of peer production through the application and evaluation of an established domain of communication scholarship around communal public goods in interactive communication systems. Third, by conducting a comparative analysis across peer production communities in order to better understand their social and organizational dynamics, we also contribute an empirical test of the prevailing assumption that peer production projects inherently advance “small-d” democracy. Finally, our results provide evidence that the widely observed hierarchies and inequalities within peer production projects are associated with enhanced oligarchic organizational structures and behavior.

Our findings show that peer production communities tend to follow Michels’ iron law

⁸<http://wikipartido.mx/>

of oligarchy, extending previous findings on the role of elites and powerful leaders in peer production and contradicting the notion the peer production and other forms of communal public goods production entail democratic organizational forms. This study therefore recalls the paradoxes of participatory processes observed by scholars of organizational democracy: there may be trade-offs and contradictions that accompany any effort to organize effective collaboration and incorporate meaningful participation at the same time (Stohl and Cheney, 2001).

We note that some wikis in our dataset appear more robustly democratic than others. Although these cases are exceptional, they signal that digital technologies, like their offline counterparts, might – or might not – be used to create participatory democratic organizations. The opportunities to do so are neither foretold nor foreclosed by the technologies themselves. Understanding why *some* peer production projects create robust democratic organizations is a promising area for future research. But this research is precluded when scholars select more democratic organizations *a priori* or simply take the participatory and democratic nature of peer production organizations for granted.

What sort of organizational culture and institutions will contemporary online associations, movements, and organizations create? The impact of peer production has grown as more organizations (firms, political parties, movements, and voluntary groups) adopt the platforms and practices of peer production to pursue their objectives. Some, such as the Occupy movement and the *Wikipartido*, have explicitly modeled their organization on wikis and other peer production communities in an effort to facilitate deeper participation and democracy. Our findings here suggest that the adoption of peer production's organizational forms may inhibit the achievement of enhanced organizational democracy.

The fact that peer production may bring with it oligarchic organizations recalls previous work on associational culture and its influence on broader social institutions. For example, when Alexis de Tocqueville visited America in the early 19th century, he saw a flowering of civic associations, the organizational structures of which he believed contributed to a democratic national culture (de Tocqueville, 2004). Similarly, when Seymour Martin Lipset and his colleagues looked to explain the remarkable democratic institutions within the International Typographical Union in the 1950s, they pointed to the robust culture of participation and engagement that cut across many levels of authority and experience within the union as a foundational factor. A few decades later, Robert Putnam bemoaned the apparent collapse of civic associationism in America, a pattern which many

have subsequently sought to connect with changes in political culture. Consistent with Michels' original concerns, the rise of labor parties in various countries of the world has also brought about the nationalization of labor politics as the organization of unions becomes part and parcel of the organization of national political culture.

Although, invoking U.S. Supreme Court Justice Louis Brandeis, online collectives have been hailed as contemporary "laboratories of democracy", our findings suggest that they may not necessarily facilitate enhanced practices of democratic engagement and organization. Indeed, our results imply that widespread efforts to appropriate online organizational tactics from peer production may facilitate the creation of entrenched oligarchies in which the self-selecting and early-adopting few assert their authority to lead in the context of movements without clearly defined institutions or boundaries.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the input, support, and feedback of Yochai Benkler, Peter Evans, Eric von Hippel, Tom Malone, Mitch Resnick, Coye Cheshire, Lance Bennett, Helen Margetts, Abigail Andrews, David Karpf, Per Selle, the MIT Economic Sociology Working Group, the Cooperation Group at the Berkman Center for Internet and Society, the Center for Information Technology and Policy, the Online Collective Action Working Group at the 2013 ECPR Joint Sessions, the Northwestern Communication Studies Faculty Research in Progress Workshop, the Buffett Center at Northwestern University, the Computational Social Science Workshop at the University of Chicago, the City University of Hong Kong, Wikimania, two anonymous reviewers, and the editor and staff of the *Journal of Communication*.

REFERENCES

- Andrews, K. T., Ganz, M., Baggetta, M., Han, H., and Lim, C. (2010). Leadership, membership, and voice: Civic associations that work. *American Journal of Sociology*, 115(4):1191–1242.
- Barakso, M. and Schaffner, B. F. (2008). Exit, voice, and interest group governance. *American Politics Research*, 36(2):186–209.
- Benkler, Y. (2002). Coase's penguin, or, linux and the nature of the firm. *Yale Law Journal*, 112(3):369.

- Benkler, Y. (2006). *The Wealth of Networks: How Social Production Transforms Markets and Freedom*. Yale University Press.
- Bennett, W. L. and Segerberg, A. (2012). The logic of connective action. *Information, Communication & Society*, 15(5):739–768.
- Bennett, W. L. and Segerberg, A. (2013). *The Logic of Connective Action: Digital Media and the Personalization of Contentious Politics*. Cambridge University Press.
- Bimber, B. A., Flanagin, A. J., and Stohl, C. (2012). *Collective action in organizations: interaction and engagement in an era of technological change*. Communication, society and politics. Cambridge University Press, New York.
- Black, L. W., Welsler, H. T., Cosley, D., and DeGroot, J. M. (2011). Self-governance through group discussion in wikipedia measuring deliberation in online groups. *Small Group Research*, 42(5):595–634.
- Butler, B. S. (2001). Membership size, communication activity, and sustainability: A resource-based model of online social structures. *Information Systems Research*, 12(4):346–362.
- Butler, B. S., Joyce, E., and Pike, J. (2008). Don't look now, but we've created a bureaucracy: the nature and roles of policies and rules in wikipedia. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '08*, pages 1101–1110, New York, NY, USA. ACM.
- Castells, M. (2012). *Networks of Outrage and Hope: Social Movements in the Internet Age*. John Wiley & Sons.
- Child, J. T. and Shumate, M. (2007). The impact of communal knowledge repositories and people-based knowledge management on perceptions of team effectiveness. *Management Communication Quarterly*, 21(1):29–54.
- Cho, H. and Lee, J.-S. (2008). Collaborative information seeking in intercultural computer-mediated communication groups testing the influence of social context using social network analysis. *Communication Research*, 35(4):548–573.
- Connolly, T. and Thorn, B. K. (1990). Discretionary databases: Theory, data, and implications. In Fulk, J. and Steinfield, C., editors, *Organizations and communication technology*, pages 219–233. Sage, Newbury Park, CA.
- de Tocqueville, A. (2004). *Democracy in America*. Library of America, New York, NY.
- Earl, J. and Kimport, K. (2011). *Digitally Enabled Social Change: Activism in the Internet Age*. MIT Press.
- Forte, A., Larco, V., and Bruckman, A. (2009). Decentralization in wikipedia governance. *J. Manage. Inf. Syst.*, 26(1):49–72.

- Fox, J. (2003). Effect displays in r for generalised linear models. *Journal of Statistical Software*, 8(15):1–27.
- Fulk, J., Flanagan, A. J., Kalman, M. E., Monge, P. R., and Ryan, T. (1996). Connective and communal public goods in interactive communication systems. *Communication Theory*, 6(1):60–87.
- Fuster Morell, M. (2010). *Governance of Online Creation Communities: Provision of infrastructure for the building of digital commons*. European University Institute, Ph.D. thesis edition. Defense date: 20/09/2010.
- Fuster Morell, M. (2012). The free culture and 15M movements in Spain: Composition, social networks and synergies. *Social Movement Studies*, 11(3-4):386–392.
- Halfaker, A., Geiger, R. S., Morgan, J. T., and Riedl, J. (2013). The rise and decline of an open collaboration system how wikipedia’s reaction to popularity is causing its decline. *American Behavioral Scientist*, 57(5):664–688.
- Halfaker, A., Kittur, A., and Riedl, J. (2011). Don’t bite the newbies: how reverts affect the quantity and quality of wikipedia work. In *Proceedings of the 7th International Symposium on Wikis and Open Collaboration, WikiSym ’11*, pages 163–172, New York, NY, USA. ACM.
- Hardin, G. (1968). The tragedy of the commons. *Science*, 162(3859):1243–1248. PMID: 17756331.
- Healy, K. and Schussman, A. (2003). The ecology of open-source software development. *Unpublished manuscript*.
- Hess, C. and Ostrom, E., editors (2011). *Understanding Knowledge as a Commons: From Theory to Practice*. The MIT Press.
- Hill, B. M., Burger, C., Jesse, J., and Bacon, J. (2008). *Official Ubuntu Book*. Prentice Hall, 3 edition.
- Ingo, H. (2006). *Open Life: The Philosophy of Open Source*. Lulu.com.
- Jenkins, J. C. (1977). Radical transformation of organizational goals. *Administrative Science Quarterly*, 22(4):568–586.
- Joyce, E., Pike, J. C., and Butler, B. S. (2012). Rules and roles vs. consensus: Self-governed deliberative mass collaboration bureaucracies. *American Behavioral Scientist*.
- Karpf, D. (2012). *The MoveOn Effect: The Unexpected Transformation of American Political Advocacy*. Oxford University Press, USA, first edition edition.
- Kittur, A., Chi, E., Suh, B., and Mytkowicz, T. (2007). Power of the few vs. wisdom of the crowd: Wikipedia and the rise of the bourgeoisie. San Jose, California.

- Kollock, P. (1999). The economies of online cooperation: Gifts and public goods in cyberspace. In Smith, M. and Kollock, P., editors, *Communities in Cyberspace*, pages 220–239. Routledge, London.
- Konieczny, P. (2009). Governance, organization, and democracy on the internet: The iron law and the evolution of wikipedia. *Sociological Forum*, 24(1):162–192.
- Kriplean, T., Beschastnikh, I., McDonald, D. W., and Golder, S. A. (2007). Community, consensus, coercion, control: cs*w or how policy mediates mass participation. In *Proceedings of the 2007 international ACM conference on Supporting group work*, GROUP '07, pages 167–176, New York, NY, USA. ACM.
- Leach, D. K. (2005). The iron law of what again? conceptualizing oligarchy across organizational forms. *Sociological Theory*, 23(3):312–337.
- Leuf, B. and Cunningham, W. (2001). *The Wiki way: quick collaboration on the Web*. Addison-Wesley Longman Publishing Co., Inc. Boston, MA, USA.
- Lipset, S. M., Trow, M. A., and Coleman, J. S. (1956). *Union democracy: The Internal Politics of the International Typographical Union*. Free Press, Glencoe, Ill.
- Loubser, M. (2010). *Organisational Mechanisms in Peer Production: the Case of Wikipedia*. DPhil, Oxford Internet Institute, Oxford University, Oxford, UK.
- Michels, R. (1915). *Political Parties: A Sociological Study of the Oligarchical Tendencies of Modern Democracy*. Hearst, New York, NY.
- Olson, M. (1965). *The Logic of Collective Action: Public Goods and the Theory of Groups*. Harvard University Press, first edition edition.
- O'Mahony, S. and Ferraro, F. (2007). The emergence of governance in an open source community. *Academy of Management Journal*, 50(5):1079–1106.
- Ortega, F. (2009). *Wikipedia: A Quantitative Analysis*. PhD dissertation, Universidad Rey Juan Carlos.
- Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, New York, NY.
- Pancier, K., Halfaker, A., and Terveen, L. (2009). Wikipedians are born, not made: a study of power editors on wikipedia. In *Proceedings of the ACM 2009 international conference on Supporting group work*, GROUP '09, pages 51–60, New York, NY, USA. ACM.
- Piven, F. F. and Cloward, R. A. (1977). *Poor people's movements: why they succeed, how they fail*. Pantheon Books, New York, 1 edition.
- Polletta, F. (2013). Participatory democracy in the new millennium. *Contemporary Sociology: A Journal of Reviews*, 42(1):40–50.

- Priedhorsky, R., Chen, J., Lam, S. T. K., Panciera, K., Terveen, L., and Riedl, J. (2007). Creating, destroying, and restoring value in wikipedia. In *Proceedings of the 2007 international ACM conference on Supporting group work*, GROUP '07, pages 259–268, New York, NY, USA. ACM.
- Raban, D. R. and Rafaeli, S. (2007). Investigating ownership and the willingness to share information online. *Computers in Human Behavior*, 23(5):2367–2382.
- Rafaeli, S. and LaRose, R. J. (1993). Electronic bulletin boards and "Public goods" explanations of collaborative mass media. *Communication Research*, 20(2):277–297.
- Rucht, D. (1999). Linking organization and mobilization: Michels's iron law of oligarchy reconsidered. *Mobilization: An International Quarterly*, 4(2):151–169.
- Schweik, C. M. and English, R. C. (2012). *Internet success: a study of open-source software commons*. MIT Press, Cambridge, Mass.
- Shaw, A. (2012). Centralized and decentralized gatekeeping in an open online collective. *Politics & Society*, 40(3):349–388.
- Shirky, C. (2008). *Here Comes Everybody: The Power of Organizing Without Organizations*. Penguin Press HC, The.
- Singer, J. D. and Willett, J. B. (2003). *Applied Longitudinal Data Analysis: Modeling Change and Event Occurrence*. Oxford University Press, USA, 1 edition.
- Stohl, C. and Cheney, G. (2001). Participatory Processes/Paradoxical practices communication and the dilemmas of organizational democracy. *Management Communication Quarterly*, 14(3):349–407.
- Tufekci, Z. and Wilson, C. (2012). Social media and the decision to participate in political protest: Observations from tahrir square. *Journal of Communication*, 62(2):363–379.
- Viégas, F. B., Wattenberg, M., McKeon, M. M., and Schuler, D. (2007). The hidden order of wikipedia. In *Online Communities and Social Computing*, HCII, pages 445–454. Springer-Verlag, Berlin, Heidelberg.
- Voss, K. and Sherman, R. (2000). Breaking the iron law of oligarchy: Union revitalization in the american labor movement. *American Journal of Sociology*, 106(2):303–349. ArticleType: research-article / Full publication date: September 2000 / Copyright 2000 The University of Chicago Press.
- Weber, S. (2004). *The Success of Open Source*. Harvard University Press.
- Wu, F., Wilkinson, D. M., and Huberman, B. A. (2009). Feedback loops of attention in peer production. In *Proceedings of the 2009 International Conference on Computational Science and Engineering - Volume 04*, CSE '09, pages 409–415, Washington, DC, USA. IEEE Computer Society.

- Yuan, Y. C., Fulk, J., and Monge, P. R. (2007). Access to information in connective and communal transactive memory systems. *Communication Research*, 34(2):131–155.
- Zhu, H., Kraut, R. E., Wang, Y.-C., and Kittur, A. (2011). Identifying shared leadership in wikipedia. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI '11, pages 3431–3434, New York, NY, USA. ACM.

APPENDIX: ROBUSTNESS CHECKS

	M4	M5	M6	M7
(Intercept)	0.135*** (0.010)	-0.259*** (0.021)	-0.002*** (0.000)	0.012** (0.004)
<i>week</i>	0.000 (0.000)	-0.001*** (0.000)	0.000*** (0.000)	0.000*** (0.000)
<i>week</i> ²	0.000** (0.000)	0.000*** (0.000)	0.000*** (0.000)	0.000*** (0.000)
$\ln accounts_{total}$	-0.002 (0.002)	0.021*** (0.005)	0.001*** (0.000)	-0.002*** (0.001)
$\ln pages_{total}$	0.001 (0.001)	0.021*** (0.004)	0.000 (0.000)	0.007*** (0.001)
$\ln admins_{total}$	-0.081*** (0.004)	0.142*** (0.009)	0.002*** (0.000)	-0.013*** (0.002)
$\ln editors_{week}$	0.086*** (0.002)	-0.102*** (0.005)	-0.001*** (0.000)	0.062*** (0.001)
$\ln edits_{week}$	-0.049*** (0.001)	0.094*** (0.003)	0.001*** (0.000)	-0.027*** (0.001)
$\ln proj-edits_{week}$	0.710*** (0.001)			0.581*** (0.001)
Num. obs.	118994	26138	112938	146858
Num. groups: wiki	554	554	683	683
Variance: 1 wiki	0.027	0.045	0.000	0.004
Variance: Residual	0.094	0.114	0.000	0.022

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table 3: Table of robustness checks. Each fitted multilevel regression model uses wiki week as its unit of analysis. *M4* is a linear model predicting the logged number of edits made by non-administrators on administrative “project” pages controlling for total edits to these pages. *M5* is a linear model predicting the proportion of edits of administrative “project” pages made by administrators. *M6* is a linear model predicting the proportion of edits by experienced editors that are reverted by administrators. *M7* is a linear model predicting the number of different editors of administrative “project” pages in a wiki week. ($n = 146,858$ wiki weeks from 683 wikis.)