

Essays on Volunteer Mobilization in Peer Production

by

Benjamin Mako Hill

B.A. Hampshire College (2003)

S.M. Media Arts and Sciences
Massachusetts Institute of Technology (2007)

Submitted to the Sloan School of Management in Partial Fulfillment of the Requirements
for the Degree of Doctor of Philosophy in Management and Media Arts and Sciences

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2013

©2013 Massachusetts Institute of Technology. All rights reserved.

Signature of Author: _____

Sloan School of Management

July 31, 2013

Certified By: _____

Eric von Hippel

T. Wilson Professor of Management and Professor of Engineering Systems

Thesis Supervisor

Accepted By: _____

Ezra W. Zuckerman Sivan

Nanyang Technological University Professor

Director, Sloan PhD Program

Essays on Volunteer Mobilization in Peer Production

by

Benjamin Mako Hill

Submitted to the Sloan School of Management in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy in Management and Media Arts and Sciences on July 31, 2013.

ABSTRACT

Although some examples of Internet-based collaborative “peer production” – like Wikipedia and Linux – build large volunteer communities and high-quality information goods, the vast majority of attempts at peer production never even attract a second contributor. This dissertation is composed of three essays that describe and test theories on the sources and effects of volunteer mobilization in peer production.

The first essay is a qualitative analysis of seven attempts to create English-language online collaborative encyclopedia projects started before January 2001, when Wikipedia was launched. Analyzing data from interviews of these Wikipedia-like projects’ initiators, along with extensive archival data, I offer a set of three propositions for why Wikipedia, similar to previous efforts and a relatively late entrant, attracted a community of hundreds of thousands while the other projects did not.

In the second essay, I use data from the Scratch online community – a large website where young people openly share and remix animations and games – to present evidence of a trade-off between “generativity” (i.e., qualities of work products likely to attract follow-on contributors) and the originality of the derivative work products that follow.

In the third essay, I consider the relationship between volunteer mobilization and governance in peer production organizations. Although large successful peer production projects have inspired scholars and social movements, I use longitudinal data of internal processes drawn from a population of wikis to show that, like other democratic organizations, peer production projects exhibit governance consistent with Robert Michels’ “Iron Law of Oligarchy.”

THESIS SUPERVISOR: Eric von Hippel

TITLE: T. Wilson Professor of Management and Professor of Engineering Systems

Thesis Committee

Yochai Benkler

Jack N. and Lillian R. Berkman Professor for Entrepreneurial Legal Studies
Faculty Co-Director, Berkman Center for Internet and Society
HARVARD LAW SCHOOL

Thomas Malone

Patrick J. McGovern Professor of Management
Director, MIT Center for Collective Intelligence,
MIT SLOAN SCHOOL OF MANAGEMENT

Mitchel Resnick

LEGO Papert Professor of Learning Research
Program in Media Arts and Sciences
MIT MEDIA LAB

Eric von Hippel (Committee Chair)

T. Wilson Professor of Management
Professor of Engineering Systems
MIT SLOAN SCHOOL OF MANAGEMENT

Contents

Contents	7
List of Figures	9
List of Tables	10
1 Almost Wikipedia:	
Eight Early Encyclopedia Projects and the Mechanisms of Collective Action	11
1.1 Introduction	11
1.2 Background	14
1.3 Failed “Wikipedias”	17
1.4 Data and Methods	20
1.5 Results	22
1.6 Alternative Explanations	37
1.7 Conclusion	41
2 The Remixing Dilemma:	
The Trade-off Between Generativity and Originality	45
2.1 Introduction	45
2.2 Background	48
2.3 Empirical Setting and Methods	52
2.4 Results	62
2.5 Limitations	64
2.6 Discussion	67
3 Laboratories of Oligarchy?	
How The Iron Law Extends to Peer Production	69
3.1 Introduction	69
3.2 Background	71
3.3 Empirical Setting	78
3.4 Data and Measures	80
3.5 Results	83
3.6 Discussion	88

References	93
Acknowledgments	103
A Almost Wikipedia Interview Protocol	109

List of Figures

1.1	Two-by-two representation of a theoretical design space in which peer production projects mobilize.	13
1.2	Two-by-two table adapted from Figure 1.1 showing placement of OCEPs.	35
1.3	Gantt chart showing when OCEPs were active.	40
2.1	Screenshot of the Scratch desktop application.	53
2.2	Screenshot of the Scratch online community front page.	54
2.3	Plots of estimated values from fitted models for prototypical Scratch projects.	64
3.1	Plots of key variables for a <i>Brickipedia</i> , an exemplary wiki in our sample.	84
3.2	Plots of estimated values from fitted models for prototypical wikis.	87

List of Tables

1.1	List and descriptive data on each OCEP started in or before January 2001. . . .	18
1.2	Table of OCEPs coded, dichotomously, for each proposition.	34
2.1	Summary statistics for variables used in analysis.	57
2.2	Summary of fitted regression models.	61
3.1	Summary statistics for all of the wikis included in our analysis.	80
3.2	Summary of fitted multilevel regression models.	86

Essay 1

Almost Wikipedia: Eight Early Encyclopedia Projects and the Mechanisms of Collective Action

1.1 INTRODUCTION

Over the last decade, several books (e.g., Reagle, 2008; Tapscott and Williams, 2008; Lih, 2009) and more than 6,200 peer-reviewed articles have been published about Wikipedia.¹ With interest driven by the rapid growth of both its encyclopedic product and community of contributors that builds it (see Ortega, 2009), a range of scholars and journalists have held Wikipedia up as a model of collaboration and collective action on the Internet (e.g., Benkler, 2006; Shirky, 2010). Although Wikipedia was not the first website to use a wiki – a piece of server software, invented by Ward Cunningham, that allows users to create and edit Web page content using any Web browser (Leuf and Cunningham, 2001) – its example has inspired countless follow-on efforts. By 2010, the wiki-hosting firm Wikia alone hosted more than two thousand websites with the phrase “pedia” in their URLs – three percent of their hosted projects.²

It would be an enormous understatement to state that not all attempts at collaboration and collective action online have enjoyed Wikipedia’s success. Indeed, research has shown

¹This number reflects the number of results from a *Google Scholar* search for articles with “Wikipedia” in the title. (Accessed June 16, 2010)

²Results of personal research using a dataset of all of Wikia’s published database “dumps” gathered in April 2010.

that few free/libre open source software projects (FLOSS) have more than one contributor (Healy and Schussman, 2003); the vast majority of wikis are similarly uncollaborative in their production (Reich et al., 2012); most remixable media never elicits a derivative (Hill and Monroy-Hernández, 2013). As a result, an important research question for scholars of peer production is: *Why do some peer production projects successfully attract contributors while most do not?* It is also one largely unaddressed in the literature on peer production to date.

Through extensive research on Wikipedia, scholars know quite a lot about *how* Wikipedia built a large community. As a group, we have documented Wikipedia’s history, organization, processes, and routines in detail. However, because we have not compared Wikipedia, for the most part, to similar failures, we still know very little about *why* Wikipedia succeeded. This paper examines seven volunteer-driven online collaborative encyclopedia projects (OCEP) founded before January 2001, when Wikipedia was launched. Although most of these projects attracted some participants, and while some built sizable groups of volunteers and substantial collections of articles, none of these projects became examples of mass collaboration on anything near the scale of Wikipedia. Using qualitative analysis of interviews from initiators of these projects and extensive archival material, I use multiple case study methods (Eisenhardt, 1989) to build support for a series of propositions that both offer a set of potential answers to the question posed by Wikipedia’s surprising success and that suggest important mechanisms of collective action in general.

Analysis of the data from initiators and archival data suggests that Wikipedia attracted a large community of contributors, while similar projects struggled, for three reasons (presented as propositions): First, (P1) Wikipedia attracted volunteer contributors because it took as its goal a model of “encyclopedia” familiar to potential contributors. Second, (P2) Wikipedia attracted volunteer contributors because it offered low barriers to contribution. Third and finally, (P3) Wikipedia attracted volunteer contributors because it deemphasized attribution and social ownership of content. Before concluding, I also consider and reject two alternative explanations: first (AE1), that Wikipedia was more technologically sophisticated than alternatives; and second (AE2) that early OCEPs were simply ahead of their

		Innovativeness of Goal/Product	
		Familiar	Novel
Innovativeness of Process/Tools	Familiar	Traditional products using traditional methods and tools. <i>"Like Encyclopedias Britannica — just online and free."</i>	New products using traditional methods and tools. <i>"A new type of encyclopedia, but produced like the old ones."</i>
	Novel	Traditional products using new methods and tools. <i>"Like Encyclopedias Britannica, but produced in a radically new way."</i>	New products using novel methods and tools. <i>"A new type of encyclopedia produced in a radically new way."</i>

Figure 1.1: Representation of a theoretical design space in which peer production projects vary between high and low levels of innovation in their goals and products (columns) and methods and processes (rows). I propose that projects in the bottom-left shaded quadrant will be most effective.

time. Taken together, my results suggest a theory of volunteer mobilization in peer production and beyond. Although Wikipedia was extremely innovative in terms of *how* the project set out to build its encyclopedia, it was profoundly uninnovative when it came to *what* it set out to build.

Early OCEPs varied in the familiarity of both their products and their processes. Figure 1.1 shows the four theoretically possible combinations of familiar (or novel) processes and products as well as providing a concise description of how an OCEP of each combination might be described. Wikipedia was the only project in my sample of OCEPs coded as falling into the bottom-left quadrant that reflects the combination of a familiar goal (e.g., “simply reproduce Encyclopedias Britannica”) with innovative methods (e.g., “anybody can edit anything”). My analysis suggests Wikipedia attracted an enormous community while very similar projects did not because Wikipedia sought to build a product that potential contributors were already deeply familiar with, while also adopting a novel set of processes, tools, and methods for organizing production. Other projects struggled both to get contributors onto the same page and to organize production in a way that took advantage of new communication technologies.

1.2 BACKGROUND

The last decade has seen a large and growing scholarly interest in online, volunteer-driven, collaborative projects. Although early interest focused on software, theories of FLOSS were extended to a variety of online communities and products under the theoretical umbrella of “peer production” (Benkler, 2002, 2006). The peer production literature has been heavily influenced by the author and software developer Eric Raymond (1999) who influentially argued that FLOSS is driven by contributions from a large number of self-motivated participants. According to Raymond, this collaborative process is inherently superior to proprietary development processes and leads to higher quality products.

Early academic work on FLOSS relied heavily on Raymond’s description to justify and frame FLOSS research. Citing Raymond’s descriptions of ad-hoc self-organization and large-scale voluntary contributions, Benkler (2002, 2006) offers a theory of peer production, using transaction-cost economics from Coase (1937) and Williamson (1981). Benkler suggests that peer production is made possible by decreases in communication costs caused by the Internet and suggests that the phenomena is unique from previously theorized organizational forms (i.e., markets and firms) by both its ability to aggregate many small contributors from individuals with diverse motivations and by its lack of traditional hierarchical and legal systems of control (Benkler, 2013). “Coase’s Penguin,” Benkler’s article coining the term peer production, foreshadows future work in that it used Wikipedia (then only one year old and largely unheard of) as an example of how theory from FLOSS could be used to explain the collaborative creation of a variety of information goods. His article also acts as a template for subsequent peer production research in its reliance on a small number of the most successful examples of peer production.

In their extensive literature review, Crowston et al. (2010) suggest that a large proportion of FLOSS research has been framed as case studies of a handful of hugely successful projects like Apache, Mozilla, and the Linux kernel (e.g. Krishnamurthy, 2005; Lakhani and von Hippel, 2003; Mockus et al., 2002; MacCormack et al., 2006, 2012). Nowhere has this been more true than in regards to Wikipedia. Wikipedia is both the most visible example of peer production and the most successful example of volunteer-based productive

mass-mobilization. As the fifth most-visited website in the world, the English version of Wikipedia has over 13 million registered contributors and more than 3.4 million articles.³ As the subject of more than 6,200 academic articles, it is also likely the most intensely studied.

Raymond's description of FLOSS has attracted scholarly criticism (e.g., Berzroukov, 1998; Krishnamurthy, 2005). Both members of the FLOSS community (e.g., Hill, 2005) and sociologists (e.g., Lin, 2005) have tried to re-frame Raymond's work as a Utopian ideal rather than a representation of empirical reality. Large cross-sectional analyses of FLOSS projects by Krishnamurthy (2002), Healy and Schussman (2003), and Schweik and English (2012) have presented empirical evidence that calls into the question claims of FLOSS collaborativeness and has shown that participation in FLOSS projects is extremely skewed: the median number contributor to a project is one (Healy and Schussman, 2003). Research on wiki communities (Reich et al., 2012; Kittur and Kraut, 2010) and remixing communities (Hill and Monroy-Hernández, 2013; Luther et al., 2010) has shown similar struggles with mobilization in the large majority of attempts at peer production.

Some researchers have adopted alternative measures of "success" that rely on downloads or the attainment of project goals (e.g., Crowston et al., 2003; English and Schweik, 2007; Schweik and English, 2012; Luther et al., 2010). Using these metrics, most successful projects are the work of a single participant or a very small group. For example, in their comprehensive study of SourceForge projects, Schweik and English (2012) show projects that have made a series of releases and created sustained commons have an average of 1.2 contributors. Among collaborative FLOSS projects, West and O'Mahony (2005) finds that many are firm-based efforts with little or no outside contributions.

Although this research shows that uncollaborativeness need not be equated with failure, uncollaborative attempts to create online commons are failures of peer production; after all, inspired by Raymond, much of the interest in Wikipedia and FLOSS stems from the benefits of collaboration; von Hippel and von Krogh's (2003) "private collective model" as-

³Statistics are taken from <https://secure.wikimedia.org/wikipedia/en/wiki/Special:Statistics> (Accessed December 8, 2010).

sumes collectivity; advantages of Murray and O'Mahony's (2007) "cumulative innovation" are irrelevant without actors reusing and recombining ideas. Peer production requires pluralistic peers to produce cooperatively. In this sense, collaboration's rarity is a critical problem for research driven by an interest in quality that, theory suggests, is grounded in collaboration.

As a result, the fundamental research question facing FLOSS becomes one largely unaddressed in the literature to date: *Why do some peer production projects successfully attract contributors while most do not?* Answering this question involves several departures from most previous research. First, research should treat the peer production *project* as the unit of analysis. Second, research should take mobilization as its dependent variable. Third, because it may be possible to mobilize individuals without eliciting contributions, research should consider the degree to which projects are able to direct the energy of participants into useful contributions toward a project's goal.

Studies of mobilization in peer production can speak to issues and scholarship on mobilization more generally. Historically, questions of mobilization have been the domain of the sociological literature on social movements. McCarthy and Zald's (1977) resource mobilization theory (RMT) uses an analogy to firms to describe movements as constituted by social movement organizations (SMOs) competing for scarce resources – usually volunteer participants or donations – within a social movement industry (SMI). That said, although RMT has played a major role in the literature on social movements, very little empirical work within RMT has considered interactions at the SMI level. Important exceptions are the work of Minkoff (e.g., 1999) and Soule and King (2008). Additionally, it has been difficult, historically, for social movement researchers to study nascent SMOs because early-stage SMOs are not visible before they are active. In her work, Minkoff relies on organizational registers that include only established, usually incorporated, SMOs. Soule and King use a dataset of New York Times reports of protests. Empirical work in the sociological literature on mobilization has never been able to capture very early failed attempts.

Because they occur entirely online, many peer production projects leave behind "digital

traces” (Wesler et al., 2008) in the form of publicly available archival data. These data are of scope and scale that was unavailable to previous scholars of mobilization. This more complete record allows researchers to explore nascent attempts at movements that did not succeed in attracting *any* additional participants or that attracted participants but did not accomplish their initial goals. The following analysis considers projects in this nascent stage. Using interviews and extensive digital trace data, I build a general, if partial, theory of early stage mobilization and suggest propositions for why some attempts at collective action succeed, while most do not.

1.3 FAILED “WIKIPEDIAS”

To construct a comprehensive sample of early OCEPs active before Wikipedia, I used characteristics cited as key qualities of the initial Wikipedia project by Reagle (2008) and others. In particular, I sought out encyclopedia projects that were English-language, collaboratively-produced, online, freely-available, and volunteer driven. A sample of projects was initially formed from consulting an existing list of online encyclopedia projects hosted on Wikipedia⁴ and through searches of news articles on the database *Lexis Nexis*. Additionally, each interviewee was asked to list any other online encyclopedia projects that they were aware of. I am confident that my current sample represents the full population of publicly announced OCEPs from the period before Wikipedia was launched.

Before January 2001, when Wikipedia was founded, there existed seven other publicly announced attempts to create English-language volunteer-driven, online collaborative encyclopedia projects. Although not all of these projects self-identified publicly as encyclopedias, each of them sought to build general purpose encyclopedias in the sense that their proposed products can be accurately described using the Oxford English Dictionary’s definition of encyclopedia: “a literary work containing extensive information on all branches of knowledge.” Moreover, each of these projects either publicly described themselves as an encyclopedia or was identified as one in the press. During interviews, initiators of these projects confirmed that their project could accurately be described as encyclopedias. Each

⁴https://secure.wikimedia.org/wikipedia/en/wiki/List_of_online_encyclopedias

Project	Total Participants	Total Articles
Interpedia	400	<50 (?)
TDEP	1 (?)	5
Everything2	50,000+ (?)	500,000+
h2g2	5,000+	13,000+
TheInfo	20 (?)	50 (?)
Nupedia	2,000+ (?)	24
GNE	300+	3-4 “test” articles
Wikipedia	500,000+	2,000,000+

Table 1.1: List of OCEPs started in and before January 2001. Details of the size of the projects in total contributors and total articles are shown. These include either the total size over the life of the project or, for the projects that continue today, the total number in December 2010.

of these projects elicited content exclusively from volunteers on the Internet. Although some made special attempts to recruit content providers, none paid for content.⁵ In this sense, each of the projects represented a form of crowd-sourcing or user generated content (Howe, 2008). Each project was published on the Internet and its content was available to the public at no cost.

Several projects operated under multiple names, as indicated below. All projects are referred to using their final names throughout this article, although quotes from interviewees or archival sources have not been altered. Each OCEP in my sample is described below in the order that they were launched.

Interpedia (also referred to as the **Internet Encyclopedia Project**) was created in late 1993 by library school student Rick Gates. The project was passed off within its first few months to a group led by Canadian technologists Douglas Pardoe Wilson and Robert Neville. The Interpedia is the only project in my sample that effectively pre-dates the World Wide Web. Interpedia was organized over email lists and USENET where its discussion group sent over 700 messages from several hundred participants. Volunteer participants were engaged in four distinct technological projects to build clients and servers. Contributions were collected for use with these clients but were never published.

⁵One project, *The Vines*, paid authors a portion of advertising revenue when their articles were viewed. Because they introduced paid labor which, theory suggests, may complicate questions of mobilization, I have not included them in the analysis.

The Distributed Encyclopedia Project (TDEP) was founded by a German computer consultant in 1997 and was re-launched in 1998 with additional content. The project did not use a dynamic web application but was designed either to host articles emailed to a group of maintainers or to link to articles on other websites from a central server. In this sense, TDEP aimed to act as an index to encyclopedia articles across the web and even to multiple articles on the same topic. TDEP never attracted any articles from users other than its initiator.

h2g2 was a project funded by *The Digital Village*, a British media and video game company connected closely to the science fiction author Douglas Adams. The goal of the site was, inspired by Adams' best-selling book series, to create *The Hitchhiker's Guide To the Galaxy: Earth Edition* – an irreverent encyclopedia about earth in the style of the fictional guide. The project was initially launched live on the BBC and was later acquired by that organization. It was spun out from the BBC in 2012 and continues to operate separately today.

The Info Network (TheInfo) was a project of the 14 years old “wunderkind” and technology entrepreneur Aaron Swartz. Before launching TheInfo, Swartz co-authored the specification for RSS – a core technology for data syndication on the web used by nearly all dynamic websites today. TheInfo attracted news coverage and an award from online entrepreneur Philip Greenspun's organization *Ars Digita* but struggled to attract contributors. It continued with a very small following for more than four years before being taken offline by a hard disk crash.

Nupedia sought to be a traditional encyclopedia project available at no cost on the web. Nupedia recruited thousands of expert participants with academic credentials and subjected articles to a rigorous review process. The project was started by Jimmy Wales and Larry Sanger – the two initiators of Wikipedia – about a year before Wikipedia was launched. Wikipedia was initially started as a side project of Nupedia and Nupedia continued after Wikipedia was launched. As Wikipedia became more successful, Nupedia was increasingly ignored and eventually abandoned.

Everything2 (E2) was founded in 1998 as a hypertext encyclopedia project created by

Internet entrepreneur Nathan Oostendorp and funded largely by the sale of the very popular technology news site Slashdot during an early period of the dotcom boom. The project was originally founded as *Everything* but was changed to E2 to reflect a major technological upgrade to the software that ran the site. E2 aimed to include more than just encyclopedic content and hosted writing in a variety of forms. The community stabilized in size and continues today with a contributor-base that is about around 5% that of Wikipedia.

GNE was started as **Gnupedia** and renamed to the recursive acronym **GNE's Not an Encyclopedia** – both references to the technology project GNU (Stallman, 2002). GNE was a project initiated by FLOSS developers inspired by an essay by Richard Stallman calling for the creation of a free encyclopedia (Stallman, 1999). The project changed its name to avoid confusion with Nupedia. Gnupedia was announced almost simultaneously with Wikipedia and initially attracted similar visibility and attention. Gnupedia continued for about a year, attracted several hundred participants who engaged in a large amount of policy and technical development, collected essentially no encyclopedic content, and eventually withered.

1.4 DATA AND METHODS

Data used in this analysis came primarily from interviews with project initiators⁶ and was supplemented with extensive archival data. All interviews were opened-ended and semi-structured. Interview protocols were designed to give interviewees space to reflect on their experience with their OCEPs. Toward this end, the protocol included broad questions about initiators' backgrounds and a series of questions about their OCEP's ideation, planning, audience, contributors, structure, and experience. At the end of each interview, I asked each initiator a series of questions about their impressions of Wikipedia and asked them to compare Wikipedia to their project. The full protocol is included as Appendix A.

With one exception, for which extensive archival data is available, I was able to contact and interview at least one representative from each early OCEP. In two cases, in-person

⁶Because the term “founder” is a contested term for several of the projects in my dataset, I use the term “initiator” to refer to any individuals involved in conceptualizing and launching an OCEP. Most, but not all, of my informants self-identify as founders.

interviews were arranged. In another case, only an email-based interview was possible. In each of the other cases, interviews were conducted over the phone or VoIP. Interviews lasted between 30 minutes and 3 hours with the median length near 90 minutes. Multiple initiators were interviewed when this was relevant and possible. All interviews – 10 in total – were recorded and fully transcribed.

Project initiators proved to be an enormously valuable source of data given my research question. In every case, informants were deeply familiar with Wikipedia, had closely followed its development, and had carefully compared their project and their own choices as initiators to Wikipedia's choices and experiences. Every single informant indicated that they considered Wikipedia a missed opportunity for themselves and their projects. As a result, each initiator had reflected, thoughtfully and at length, over a span of more than a decade, on the core research question at the heart of this study.

Additionally, I assembled extensive archival data for each project. I downloaded full email or USENET archives for Nupedia, GNE, and Interpedia which document both the planning stages and the period after "launch." Similar material was not created or was not available for other projects. Using archival information and discussion forums available on the web and taken from the Internet Archive's *Way Back Machine*, I gathered documentation of process discussions for E2, TheInfo, and h2g2. Additionally, I acquired a subset of relevant email from the archives of a major advocate of online encyclopedia projects with whom several of the projects had corresponded in depth during the planning phases of their projects. Finally, I did full text keywords searches for each project in Lexis Nexus and assembled lists of published news articles for each project from before 2001. The total archival dataset was more than 3,000 pages of text.

Data analysis was structured as a comparative multiple case study (see Eisenhardt, 1989). In most cases, I was able to confirm both descriptions and ascriptions by project initiators using available archival data with statements by the initiators of other projects, and, where multiple initiators were interviewed, by statements of co-initiators. All interviews were hand coded through an iterative step-wise process using the RQDA qualitative data analysis software (Huang, 2012). Using methods from grounded theory (Charmaz, 2006),

codes were initially inductive and emergent from the text. Subsequently, data was also coded using deductive codes representing concepts from theory (see Corbin and Strauss, 2007). As new codes were created or merged, data was reviewed and recoded in an iterative, step-wise manner. Codes were categorized and groups were combined into higher-level meta-codes to create themes.

1.5 RESULTS

The process of coding and code aggregation led to three distinct themes associated with propositions for why Wikipedia succeeded. Each of these propositions was raised by initiators from at least projects and these propositions reflect the full list of themes cited by at least three projects as explanations for Wikipedia's success. These propositions are:

P1: Wikipedia attracted contributors because it was built around a *familiar product*.

P2: Wikipedia attracted contributors because it offered *low barriers to contribution*.

P3: Wikipedia attracted contributors because it offered low attribution and *low social ownership of content*.

Each of these propositions is discussed in depth in the subsections below.

Proposition 1: Familiar Product

Many respondents suggested that Wikipedia was effective in attracting contributors because it took steps to clearly define itself as an encyclopedia. On the other hand, there was considerable variation in how closely other OCEPs hewed to traditional conceptions of “encyclopedia-ness” and many projects attempted to extend the traditional encyclopedia genre to explore new possibilities for encyclopedic reference works online. Wikipedia has called itself an encyclopedia on every page since its establishment. Through this type of appeal to tradition, Wikipedia was among the most conservative projects in my sample in terms of its goal. Despite the new forms of reference works made possible by the Internet, Wikipedia set out to be just an encyclopedia.

In his book and in several articles, media theorist Joseph Reagle places Wikipedia into historical context by describing it as the latest in a long series of attempts to codify and transfer knowledge through reference works that he traces back to Diderot and the Enlightenment (Reagle, 2009, 2008; Loveland and Reagle, 2013). Reagle suggests that Wikipedia's core policies – e.g., neutral voice, coverage of only notable subjects, a prohibition on original research, authority built through references to external published sources – can all be traced to an understanding, familiar to potential contributors, of what an encyclopedia should be and has been.

Many other OCEPs were more ambitious in the way they defined their projects' goals. By offering subtly different ideas of *what* they were building, other OCEPs complicated and elaborated on the idea of an encyclopedia in describing their work. For example, in one of the very first articles shared on his site, E2 initiator Nate Oostendorp explained:

Everything is a flexible web database, created by [Block Stackers Inc.] which seeks to find the best way to store and link ideas. The result: it's absolutely crazy... (Oostendorp, 1999)

E2 struggled with explaining to users what it sought to be without constraining itself to existing genres. That said, Oostendorp's answer may not have been of much guidance to prospective E2 contributors trying to decide what to contribute. An early E2 "frequently asked questions" (FAQ) document included, as its very first entry, the question, "What is Everything?" In answering the question, E2's only paid editor, struggled with the difficulty of explaining to new users exactly what E2 was and was not:

Everything is what you make of it, that's the bottom line. It's open-ended, open-minded and waiting for you. You can node your diary, CD collection, dreamlog, notes on the apocalypse or a recipe for fettuccine alfredo. You can sit around and read what other people have written. You can recommend changes in the system. You can do almost anything you want to [providing] you spell it right...

Everything is an ever-growing, pulsating database that moves through cyberspace like a death-borg...slowly collecting and assimilating information and nonsense until...until... (Postma, 2000)

In interviews, multiple early E2 contributors confirmed that despite the fact that a large proportion of early content in E2 was encyclopedic in nature, users had trouble learning *what* to contribute. As a result, while the site was able to build a following of thousands of contributors, it failed to take off in the way that Wikipedia did.

In an interview, Aaron Swartz, initiator of TheInfo, recalled that his project similarly resisted a strict encyclopedic frame:

I don't think I conceived of it as like "let's just put an encyclopedia online." I think I probably [thought] like, "this is going to be an exploration and we're going to figure out what a reference work online looks like" (Interview).

Elsewhere, Swartz explained that he was highly influenced by writers like Vannevar Bush and Theodore Nelson who each described a science fiction-like future of interlinked databases that went beyond what was possible in traditional encyclopedias (see Bush, 1945; Nelson, 1981). Swartz, like other initiators, saw his project as an opportunity to explore these possibilities and sought to do so by not limiting his project to existing paper-based conceptions of what a reference work should be.

Although h2g2 was more reserved in its ambitions, it reflects yet another example of an OCEP that departed from the traditional encyclopedic frame. In Douglas Adams' books, the *Hitchhikers Guide* is an electronic encyclopedia that forms a compendium of essential knowledge about the universe and that is described as being humorous and irreverent in tone – especially when compared to stodgier *Encyclopedia Galactica* which is described by Adams as a traditional encyclopedia. Avoiding the term encyclopedia, h2g2 explained at its launch that it was, "a repository of human experience and knowledge" (Archival Data). As was typical of other OCEPs, h2g2's minor deviation from the well-understood encyclopedia model caused friction with contributors who struggled to understand what

it was they were being asked to contribute. One h2g2 initiator describes this difficulty in a way that was echoed by other interviewees:

So one of the problems was firstly that people would be writing completely fictional stuff about the universe. Y'know, about the Hitchhiker's universe. And we'd go, "No, no, no, no. You're not getting it. This is for real people. This is about the real world."

Then what they did, at the same time, was that they'd also do stuff about the real world, but try and write it from the point of view for an intergalactic guide. So we'd get articles about soccer that would start with, "On the planet earth which is the third planet out from the solar system Sol, the humans like to play, blah blah blah blah." Shut up. Alright? It's like, "this is going to be read by humans who live on earth."

We had piles and piles of that shit. And we had to shovel our way out from under it (Interview).

Even though h2g2 was encouraging the production of encyclopedia articles, the fact that they described their project as the *Earth Edition of The Hitchhiker's Guide to the Galaxy* confused potential contributors. Although most potential contributors to OCEPs arrived with some familiarity with encyclopedias, every h2g2 contributor first had to learn how h2g2 was different. The quotation above, and many others like it, suggest that even when OCEPs deviated only slightly from a traditional encyclopedia in their stated goal, this deviance significantly hindered mobilization.

Like many other interviewees, E2 initiator Nate Oostendorp suggested that E2's attempt to avoid a pure encyclopedic framing was the major reason that E2 failed to mobilize contributors to the degree that Wikipedia did:

I don't think we ever used the term encyclopedia and that probably would have been smart...

Wikipedia had a much more focused purpose than Everything2. Everything2 was, just by its nature, sort of zen koan like, everyone who was involved with it thought it completely defined description. And that, I think, was to its, ultimately to its detriment.

Versus Wikipedia which was like, "we're going to be like the encyclopedia, like the World Book Encyclopedia but huge and comprehensive. But we're going to keep this impartial tone and everything has to be referenced and that sort of thing" (Interview).

We can think of this contrast between models in terms of Goffman's (1974) theory of frames. The social movements literature has relied heavily on the conception of "mobilization frames" (see Snow et al., 1986; Benford and Snow, 2000). With a strong focus on frame "resonance," collective action framing theory suggests that social movements' major problem in mobilization is the construction of frames that are likely to motivate people to join or to continue participation. That said, the mobilization framing literature tends to treat goals as constant, and to focus on *how* work should be done (e.g., through discussion of a "revolutionary frame"), rather than on the role that frames play as goals or targets.

The broader sociological literature provides us with several other tools to think about how clear definitions of genres might impact the success of organizations. Zuckerman (1999) uses firm performance data to argue that a lack of conformity to categories – measured by under-coverage by stock analysts – can result in undervaluation of publicly traded firms. In the innovation literature, Hargadon and Douglas (2001) suggests that new innovations face barriers to diffusion when users do not understand how to use them. Hargadon and Douglas describe how Edison was able to sell electrification only after he re-framed it as a safer form of gas lighting through steps that included reducing the brightness of lights to more closely match what consumers expected from gas.

Work by scholars in information systems and organizations have explored another category-based effect in the adoption and effective use of new technological products. Building on Goffman (1974), Orlikowski (1992) and Orlikowski and Gash (1994) show

that collaborative software might be adopted, but used with only a subset of its features or abilities, if users do not approach the new tool with a cognitive frame that can make sense of the new functionality. In order to take advantage of new modes of work within an organization, users must understand what that new type of work could entail. In this research, users' inability to frame the nature of the processes at play is associated with the failure of new technologies. Although they use a different success metric, Schweik and English (2012) find that the strongest predictor of early-stage success for FLOSS projects is the number of categories that projects are labeled with in their database. Schweik and English are surprised by this result and attribute it, tentatively, to strong leadership. An alternative explanation is the number of categories reflects a framing effect of the type I propose.

Proposition 2: Low Barriers to Contribution

Many theories of Wikipedia's success have emphasized the low barriers associated with contribution. Benkler's (2006) theory of peer production is based, fundamentally, around the idea that Internet-mediated communication has lowered transaction costs enormously. This model suggests that many people have a latent motivation to participate in the provision of public goods on the Internet that is higher than the direct costs of doing so, but that this motivation is frequently still too low to offset barriers in the form of transaction costs associated with contributing through contracting on a market or creating or joining a firm.

For example, if a hypothetical user finds a small error while reading an encyclopedia article – e.g., a missing comma – she may be motivated to spend a small amount of time and effort to fix the error and share that modification. That said, in order to contribute, her motivation would have to be high enough to offset the costs of making that change. If the user is reading Encyclopedia Britannica without the Internet, making the change would involve finding the address of the publisher, finding a letter and envelope, writing a letter notifying the editors of the missing comma and noting the page and edition, paying for a stamp, and sending the letter. The entire experience might take half an hour and

cost US\$1 in materials. Although many people might notice the error and be motivated to fix it, fewer users would be motivated enough to invest the time and money to follow through. If, on the other hand, there is an *Edit* button at the top of the page that the user is reading, the act of fixing the comma for others may have a cost much closer to zero, and the comma will be fixed more often.

Alluding to this dynamic, many OCEP initiators suggested that Wikipedia mobilized large numbers of volunteers because their OCEPs introduced what, with foresight, they consider relatively unimportant barriers to contribution. Several initiators suggested that Wikipedia succeeded because it made contributing easy – especially for lightly motivated potential contributors. Frequently cited examples of barriers to contribution included account creation, the requirement to log in to contribute, the need to learn markup or simple codes to format text and to create hyperlinks, and the requirement to preview changes before submitting them.

For example, an initiator of h2g2 explained that he felt that Wikipedia was easier to edit than his project in part because it did not require an account or any longer-term engagement with the project:

It's... I think definitely the immediacy of it. And certainly one of the aspects was the fact that you don't have to sign up to edit. That you can look at a page and see something wrong and immediately edit without having to do anything else. Y'know? You can come along and do a drive-by-edit and never be involved again, and make a contribution. You can't do a drive-by on almost any other project (Interview).

Many initiatives voiced a similar sentiment. An initiator of GNE suggested that Wikipedia had succeeded where GNE failed because, “Wikipedia was so much simpler” (Interview) and explained that – both socially as well as technologically – Wikipedia did more to reduce barriers to contribution.

Several initiators referred to the fact that contributing to Wikipedia required relatively little or no association with the project and pointed to the absence of a steep “learning

curve” which might hinder joining. Indeed, the idea that even relatively small barriers can prevent the adoption of a technology-based product is a fundamental assumption behind the academic and practitioner focused literatures on usability (e.g., Nielsen, 1993; Preece, 2000). Several project initiators suggested that contributions by previously uninvolved contributors to Wikipedia were easier relative to other OCEPs.

While explaining why his project failed to take off while Wikipedia did not, the initiator of TDEP suggested that costs associated with contributing to the project were the major reason that TDEP failed where Wikipedia succeeded only several years later:

Well, the distributed encyclopedia failed ... Because of building encyclopedia articles using hand-crafted HTML still being too complex. Wikis solved the latter issue nicely (Email Correspondence).

In the case of TDEP, barriers included learning technology like the HyperText Markup Language (HTML) in order to contribute. Although HTML is widely known, and although contributing to Wikipedia also requires learning its own idiosyncratic markup, Wikipedia’s markup (called “wiki text”) offers a lower barrier to contribution because it looks less like “code” and more like plain text. As a result, people that do not know wiki-text can still make workable, if imperfect, contributions (Leuf and Cunningham, 2001).

On the other hand, many initiators expressed skepticism that barriers to contribution reflected a full explanation for Wikipedia’s success. For example, only one initiator (TDEP) cited barriers to contribution as the single most important reason for Wikipedia’s success relative to their OCEP. Additionally, several initiators argued that their OCEP was effectively as easy to contribute to as Wikipedia. For example, Aaron Swartz pointed to several small barriers associated with contributing to TheInfo:

So, one problem was a mandatory preview step before you saved it. Which probably wasn’t enough to kill the site single handedly, but I probably would have changed it [if given the opportunity to the project again]. Wikipedia does fine without that (Interview).

Swartz, like several other initiators, suggested that Wikipedia's success without certain barriers might have altered their own designs. But many felt that Wikipedia's barriers were still significant and that learning to use a computer, being comfortable with wiki text, and navigating Wikipedia's complex social organization, were impediments on the level of barriers in their systems. To this day, the Wikimedia foundation cites usability issues as a major barrier to contribution in Wikipedia.

Although initiators were able to point to barriers to contribution in their systems that were higher than in Wikipedia, they also pointed to many technical similarities and to relatively lower barriers in their OCEPs. Wikipedia launched with off-the-shelf software that imposed many limitations relative to other OCEPs. Finally, and although the determination is inherently subjective, I felt that, in 2001, several OCEPs appeared and operated more like Wikipedia does today than Wikipedia did when it launched.

Proposition 3: Low Social Ownership

In the literature on peer production, reputation is frequently cited as an incentive for contribution (Lerner and Tirole, 2002; Benkler, 2006, 2013). This theory suggests that individuals are paid in reputation or status in the absence of monetary payments. Toward this end, many creators of peer production and social computing systems have designed systems that provide contributors with reputation and status through identity systems that attribute contributions to their authors (Monroy-Hernández et al., 2011; Kraut and Resnick, 2012). Although initiators of OCEPs often took these ideas for granted in the design of their own systems, several pointed to the fact that Wikipedia's own design and technology provided very little indication of authorship. This, they suggested, led to low degrees of territoriality (Brown et al., 2005) and social ownership of content. This, initiators argued, helped Wikipedia mobilize contributors relative to their efforts.

Early work on incentives in peer production was influenced by economics and treated FLOSS's ability to mobilize large communities in the absence of monetary rewards to be confounding. Several early FLOSS studies looked to reputation as a form of alternative currency. For example, Lerner and Tirole (2002) suggested that people contribute to FLOSS

because they can build their reputations by doing so and they will be able to translate this reputation into higher human capital and higher paying jobs. Reputation can also be internal to communities where one's contributions to a public good are rewarded through the increased opinion, and subsequent actions, of other members of the community (Willer, 2009b,a). This approach to reputation systems has been cited more broadly as an important driver of contributions in online communities and social computing systems (Dellarocas, 2006, 2010). For examples, FLOSS projects hosted on SourceForge have administrators, developers, junior developers, and other categories of contributors.

A pre-requisite to reputation building, of course, is a way for users to attach their identity to their contributions. Many FLOSS systems reflect reputation through formal roles within projects. However, when one visits a Wikipedia article, it is unclear who has produced it (Viégas et al., 2007). Several OCEP initiators cited Wikipedias' absence of a system for easily determining who did what as a reason for Wikipedia' success. Counterintuitively, they argue that Wikipedia attracted contributors because low attribution facilitated less individual social ownership of work products and less socially risky collaboration. Most failed projects used stronger attribution which led to more "territoriality" (Brown et al., 2005). Many OCEPs allowed little or no direct collaboration on text (i.e., collaborators would need to copy and replace text to improve it). Most of the projects that did allow some direct collaboration still had explicit authorship norms (see Jaszi, 1991; Coombe, 1998) which identified an individual or group as the "author" of any particular piece of content.

The Interpedia was one project that did not allow for collaboration at the level of article text. Interpedia's design called for individuals to work together to build an encyclopedia, but to have each article produced individually. For example, one initiator of Interpedia suggested:

People would contribute articles and then they would be missing or they'd be somewhere down the list below the default article. But they'd be the work of one identifiable person. Unlike the Wikipedia thing where it'd be very difficult

to track down who contributed some messages (Interview).

A different initiator of Interpedia made a similar point:

But the Wikipedia sort of conquered because anyone could just write anything on any page without anyone's approval (Interview).

In this way, the design of several OCEPs called for writing articles that were the identifiable work of a single individual. Of course, all content in all OCEPs was licensed permissively so that anyone could modify and build on any piece of contributed content. But for projects like Interpedia and TDEP, contributors were expected to ask the original author of each article to update the text, or to write their own version of the article by incorporating their work into a significantly different article to which they could claim authorship.

Another project that did not allow direct collaboration on text was E2. In E2, each topic or "node" could have a number of different "write-ups" or articles associated with it. Each of these write-ups was listed as being produced by the author who started the write-up. Although it was possible for some users to edit others' articles, the fact that each write-up was labeled as having been written by an identifiable author led to a social norm that looked down upon this type of intra-textual collaboration. If an E2 user saw something wrong with an article, they were encouraged to provide an improved write-up of their own rather than trying to edit somebody else's. E2 initiator Nate Oostendorp explained that the clear presence of authorship led to failures to mobilize by comparing their design to Wikipedia's:

In Wikipedia, when you submit content, you don't really get an authorship credit directly. Y'know, you appear in the history but these aren't necessarily your words, they are just sort of your contribution to Wikipedia. But with Everything ... their writings were still theirs. They had control of them on the site. And they received direct attribution. I think there was some weakness there in that when people wrote something and if it was factual content and if

they had information that was incorrect, there was no real... I mean, occasionally an editor would go in and change the content but otherwise it was sort of up to them to receive communication and re-add to it (Interview).

Indeed, authorship and credit were cited repeatedly by initiators as more important than technical workflow limitations. Another early contributor and moderator of E2, Cliff Lampe, expounded on this observation:

So I think having one article as opposed to several write-ups on a node took advantage of marginal contributions in a way that E2 is not set up to. That really helped make it a much more strongly “many hands make lighter work” type of exercise (Interview).

In the projects in my sample, ownership over text was often purely social. Several OCEPs whose initiators suggested that textual ownership hindered contribution allowed, from a technical perspective, for at least some direct collaboration on the text. But because they also listed authorship explicitly, a sense of social ownership over text meant this technical possibility was rarely taken advantage of.

Of course, Wikipedia is not immune to social ownership of text. For example, Thom-Santelli et al. (2009) have described territoriality in Wikipedia. And although Thom-Santelli et al. suggest that territoriality has benefits, they point to drawbacks within Wikipedia as well. By placing a name on the top of the page, many OCEPs made it socially risky to engage in direct collaboration on text. Although counterintuitive, this result finds some support in recent research in social computing (Bernstein et al., 2011); in some cases, users can engage and collaborate *because* there are not clearly identifiable “toes” to be stepped on. When these toes are visible, many potential collaborators chose not to build on each others’ work. Ironically, the fact that Wikipedia made authorship less scrutable opened the door to deeper and more widespread collaboration.

Project	P1: Familiar Goal	P2: Low Barriers	P3: Low Social Ownership
Interpedia	No	No	No
TDEP	Yes	No	No
GNE	Yes	No	No
Everything2	No	Yes	No
h2g2	No	Yes	No
Nupedia	Yes	No	No
TheInfo	No	Yes	Yes
Wikipedia	Yes	Yes	Yes

Table 1.2: Dichotomous codes for each encyclopedia project for each of the three propositions described in the result. A code of “Yes” suggests that there was strong support for that theme in the data associated with the project while “No” suggests there was not strong support.

Synthesis

Results of data analysis suggest three proposition to explain Wikipedia’s relative success at mobilizing volunteers: (P1) a reliance on a product or goal familiar to potential contributors; (P2) low barriers to contribution; and (P3) low social ownership of content. These results are summarized for each project in dichotomous codes shown in in Table 1.2. Several patterns can be seen in these results. First, Wikipedia is the only project coded as having satisfied each proposition. Second, P2 and P3 are largely coincident in the dataset – i.e., only two projects (E2 and h2g2) have different codes for P2 and P3. And although they are conservatively coded “No” for P3, both h2g2 and Everything *did* allow some editing of each others content. It is also the case that no project other than Wikipedia is coded “Yes” for P1 is also coded “Yes” for *either* P2 or P3.

This pattern of results can be understood more clearly if the propositions are compared with traditional encyclopedias. Toward that end, we can return to the diagram in Figure 1.1. Figure 1.2 reproduces that diagram but omits descriptions of each cell and, instead, uses the codes in Table 1.2 to place each OCEP onto the grid. In the figure, P1 is mapped onto the x-axis and reflects the novelty of the product or goal. Projects in the left column are the OCEPs that hewed closely to existing definitions of encyclopedias in terms of their goal and scope. In the right column are projects that tried to expand upon this concept. In the

		Innovativeness of Goal/Product	
		Familiar (P1=Y)	Novel (P1=N)
Innovativeness of Process/Tools Novel (P2,3=Y) Familiar (P2,3=N)	Familiar (P2,3=N)	TDEP GNE Nupedia	Interpedia Everything2 h2g2
	Novel (P2,3=Y)	Wikipedia	The Info Network

Figure 1.2: Two-by-two table adapted from Figure 1.1. Propositions are added to the labels on the axes and the names of OCEPs are placed onto the grid based on their coding as described in Table 1.2.

context of traditional encyclopedias, P2 and P3 each reflects innovation in the process and tools used to organize production. This reflects, for example, what Benkler (2006) suggests is Wikipedia’s, “radically new form of encyclopedia writing.” Traditional encyclopedia writing was marked by intentional barriers to participation by non-experts and by rigid systems of workflow, roles, and processes designed to ensure quality. Because P2 and P3 were largely coincident in the dataset, I have mapped them, in combination, to the y-axis in the Figure 1.2. Projects in the top row aimed to organize production, authority, and social ownership in ways that tracked traditional methods. Projects in the bottom row used innovative processes and sought to remove barriers to contribution and social ownership.

TDEP, GNE, and Nupedia are coded as residing in the top-left quadrant. These three projects sought to build traditional encyclopedias using traditional processes and methods. They differed from more traditional encyclopedias like Britannica primarily in that they released their content freely on the Internet and did not pay authors. Each of these projects struggled to mobilize contributors and to compete with existing encyclopedias. In the bottom-right quadrant is TheInfo which reflects the only project coded as attempting to expand upon traditional encyclopedic frames while discarding traditional methods of production. Despite press and attention, TheInfo struggled to attract a community.

Projects along this diagonal (top-left to bottom-right) were the least successful efforts in the sample in terms of their ability to mobilize contributors. In general, efforts that used a mix of innovation and familiarity worked better than either very innovative or very familiar projects.

Projects on the other diagonal (top-right to bottom-left) are coded as mixing innovation with tradition. Would be contributors to projects on this diagonal were likely to be comfortable and familiar with either the goal or the process at the core of the project, but not with both. Three of the four projects along this diagonal are the projects that built sustainable communities and persist today. Interpedia is coded as falling into the top left quadrant. E2 and h2g2 are coded as having low barriers to entry but high level of social ownership over work products. Because of their mixed coding on P2 and P3, they are placed on the border of the two cells in the right column. These three projects sought to build innovative genres of reference works, but did so using methods that were at least moderately similar to traditional models. In the bottom-left shaded corner is Wikipedia, the only projects that sought to build a traditional encyclopedia using novel methods.

Taken together, these propositions suggest that Wikipedia succeeded in building a large community of volunteer contributors because it sought to build a product that potential contributors were already familiar with, but aimed to do so in a way that used novel peer productive forms of organizing volunteer labor that took advantage of advances in communication technology. Although Wikipedia's initiators encouraged editors to cast aside preconceptions of *how* an encyclopedia should be written, they appealed explicitly to their contributors' existing conceptions of what an encyclopedia should be. My results suggest that Wikipedia succeeded because it sought to do the old thing in a new way.

Consideration of other peer production projects can find anecdotal support for this dynamic. Many of the largest and most successful peer production communities would, if placed on the grid, fall unquestionably into the shaded quadrant. GNU/Linux, frequently referred to simply as the Linux operating system, began as a project to reproduce UNIX. Even the name "GNU" stands for "GNU's Not Unix" – a joke because GNU's not UNIX in name only. The popular OpenOffice and LibreOffice office suites are free software re-

implementations of Microsoft Office. The GCC compiler and the Apache webserver each implement published standards where the project's stated goal was literally written down in a standards document before any contributor arrived. OpenStreetMap seeks to create a free replacement to Google Maps.

1.6 ALTERNATIVE EXPLANATIONS

Although the inductive approach used in this paper can produce new theory, it cannot test these theories and it cannot reject the possibility that there are other valid explanations that did not emerge from my data through my methods. That said, there are two commonly cited explanations for Wikipedia's success that found disconfirming evidence in my data and analysis. First, there is little support in my dataset for the argument that (AE1) Wikipedia succeeded because it was technologically superior to alternatives. On the contrary, several project initiators cited the relative technological unsophistication of Wikipedia as an indirect reason for its success. Second, there was little support for the argument that (AE2) Wikipedia succeeded where other projects fell short because Wikipedia's timing was correct and that other projects were simply ahead of their time.

Alternative Explanation 1: Technological Superiority

If not a technological advance itself, Wikipedia is frequently described as the product of a technological inevitability. Evidence in my data set provides little support for the argument that Wikipedia succeeded because of its technological superiority. By definition, every OCEP used the Internet. Indeed, most OCEPs in my sample used technology at least as sophisticated as Wikipedia's and most used technology that was more sophisticated. For example, an initiator of Interpedia – a project that predated Wikipedia by nearly a decade – explained:

A lot of the stuff in the Wikipedia is extremely obvious and not very sophisticated. I mean, the Wikipedia is not high tech. I always imagined something high tech. That's my nature. I envision things that are of a higher technical level. We envisioned for the Interpedia as something that would be high tech.

And we could see the Interpedia inspiring the Wikipedia, but not the other way around (Interview).

On the other hand, many other projects had a strong focus on technological development. For example, TheInfo's initiator Aaron Swartz explained he saw his own role as a technological facilitator saying:

I had this notion that my job was to provide the platform ... And so, I kept trying to refine the user interface and things like that to make it more inviting so more people would write stuff (Interview).

Similarly, contributors to GNE worked on code over many months and produced several competing technological implementations of software to run the system that each went through multiple releases.

On the other hand, Wikipedia's initiators were, by far, the least technically sophisticated founding team in my sample. Every other initiator, with the exception of Nupedia which shared a set of founders with Wikipedia, was a technologist. In each case, initiators thought of their projects in primarily technological terms. Despite the fact that none of the questions in my interview protocol explicitly mentioned technology or software, a large portion of many interviews focused on issues of the design of the technological systems that facilitated interaction.

Wikipedia is the one OCEP in my sample that did not begin by writing its own software. Instead, it used an off-the-shelf, freely available, piece of wiki software called *UseModWiki*.⁷ As non-technologists initially more focused on Nupedia, Wikipedia's founders invested very few resources into technology for the project. Two OCEP initiators suggested that Wikipedia may have succeeded due to this difference in focus and their own myopic focus on technology. While a strong technological orientation was rarely cited as hurting OCEPs, it may have distracted from the key social and organizational problems of building a community of contributors.

⁷A good resource with details on UseModWiki is the Wikipedia article on the subject: <http://en.wikipedia.org/wiki/UseModWiki> (Accessed June 29, 2013)

Alternative Explanation 2: Timing

Another commonly cited alternative explanation for the success of Wikipedia relative to other OCEPs is that Wikipedia simply had better timing. Each of the seven projects in my sample was founded before Wikipedia. As a result, one explanation for Wikipedia's success was that building an OCEP was an idea whose time came in early 2001 and that earlier projects were simply ahead of the curve. Had they been started when Wikipedia was, this alternative explanation suggests, other OCEPs might have succeeded instead. Support for this perspective come from the fact that Internet use was expanding rapidly around 2001 with millions of new Internet users being added each month.⁸ Additionally, Internet technology was being developed rapidly around the turn of the millennium. OCEPs founded before 2001 may have simply been at a disadvantage relatively to Wikipedia in that they were mobilizing from a smaller pool of potential contributors and using less sophisticated tools. Although I found some support for this explanation in one case, my data suggests that timing was not a critical factor in explaining the relative failure of most OCEPs in my sample.

Interpedia was launched in 1993 when the web was still in its infancy. One of the initiators of Interpedia reflected on Interpedia's failure saying:

It could have been timing. That was part of it. We didn't have the mechanism of the web which is clearly a lot better way. Not just hypertext, but the web, as it's become so large. And that's another thing. We didn't have the installed user-base – for want of a better term – of people that have time. And so, [for Wikipedia], the timing was right (Interview).

Citing both the state of technological development and the state of potential contributor-bases, both Interpedia initiators that I interviewed alluded to timing as a possible contrib-

⁸The International Telecommunications Union publishes data on Internet penetration rates globally. Details and downloadable data are available online at <http://www.itu.int/en/ITU-D/Statistics/Pages/stat/default.aspx> (Accessed June 29, 2013). However, there is very little good data available on Internet penetration in the mid 1990s. Data from the Internet Systems Consortium suggests that connected hosts increased from several hundred thousand hosts to several tens of millions – an increase of over 100 times over the 1990s: <https://oldmirage.isc.org/solutions/survey/history> (Accessed June 29, 2013).

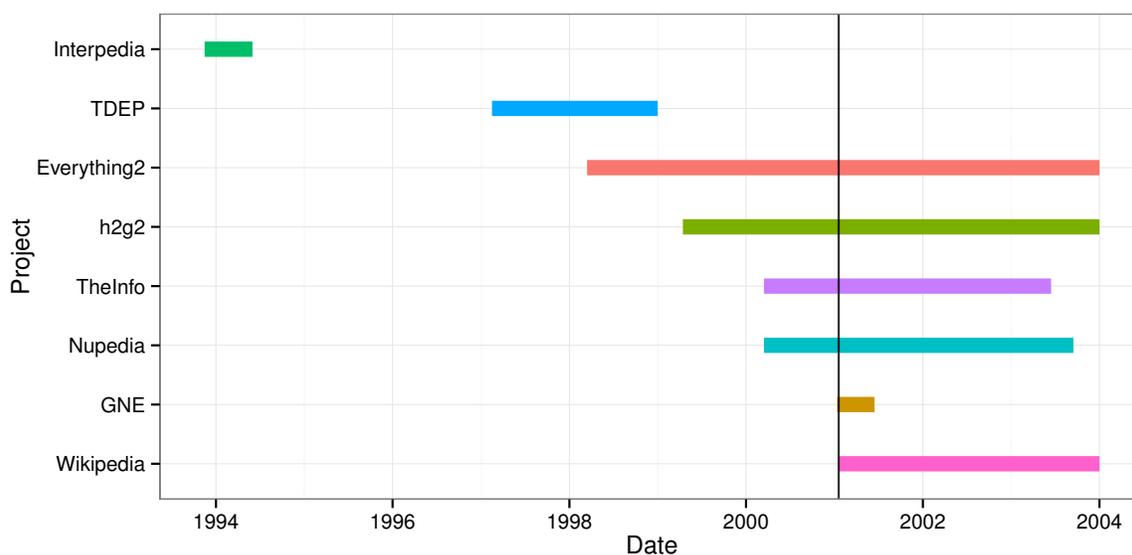


Figure 1.3: Gantt chart showing when OCEPs were active. Everything2, h2g2, and Wikipedia remain active today. The vertical line marks January 18, 2001 – the date that Wikipedia was publicly launched.

utor to their own projects struggles to mobilize relative to Wikipedia. Indeed, Interpedia was founded nearly one decade before Wikipedia at a point when the Internet was much smaller than it was in 2001.

But Interpedia was the only project that mentioned timing as a reason for their effort’s failure to mobilize volunteers. All of the other OCEPs in my sample launched within 4 years of Wikipedia’s founding and were active less than two years before Wikipedia launched. This includes TDEP which was the only other project that ended before Wikipedia launched. With the exception of Interpedia and TDEP, every project was launched within three years of Wikipedia; three were launched less than one year before, and one project launched less than a week before Wikipedia in January 2001. A plot of the periods during which each project was active is shown in Figure 1.3. In five of the seven early OCEPs, efforts continued in the period before, during, and after Wikipedia’s launch. Moreover, Wikipedia grew slowly over its first year. Several of the OCEPs in my sample – most notably E2 and h2g2 – were large and more established communities into 2002 or 2003.

It was also not the case that Wikipedia made use of technology unavailable to other projects. As detailed in Section AE1 above, Wikipedia launched using off-the-shelf software, available under a free software license. The software Wikipedia used was released in 1999 and was written in Perl – the same programming language used in several of the other OCEPs. As a result, Wikipedia’s exact technology would have been easy to use, customize, and build upon for other OCEP initiators. Indeed, archival data from TheInfo shows that a contributor suggested that the site administrators consider using wiki software saying, “there’s a thing called WikiWiki that you should go check out” (Archival Data). Initiator Aaron Swartz rejected the suggestion, preferring the technological affordances of his own system, explaining later that, “I was frustrated with [wiki software]. And you can see that I fixed a lot of things that ended up being in [later version of] Wikipedia” (Interview). Although timing may have played a role in the earliest OCEPs in my sample, it was rejected as an explanation by most projects initiators and found little support in my analysis.

1.7 CONCLUSION

Previous scholarship has attributed Wikipedia’s success in mobilizing to an accident or simply to luck (Reagle, 2009). But what is luck? Luck, according to one perspective, is simply a term we invoke when we do not understand and cannot control the processes that determine some fortunate outcome. Through this work, I have attempted to deconstruct and explain Wikipedia’s “luck” to offer a theory of volunteer mobilization in peer production. I have suggested three propositions for why Wikipedia succeeded in mobilizing a large community of volunteer contributors where very similar efforts failed. I have suggested that Wikipedia succeeded because it sought to build a product that was familiar to potential contributors while adopting innovative processes and methods of organizing labor. This pattern, I suggest, is common to many of the most successful examples of peer production.

There are limitations to the methodology employed in this essay and my results are only propositions. In future work, I hope to operationalize these propositions and test them as hypotheses in quantitative analyses. There are several large databases of FLOSS

projects like SourceForge, large datasets of populations of wikis (e.g., Kittur and Kraut, 2010; Reich et al., 2012), and remixing communities (e.g., Hill and Monroy-Hernández, 2013), which, if properly coded in terms of these propositions, may support empirical tests of the relationships suggested.

But if luck refers to a lack of control, we must also acknowledge that even successful tests of these propositions may not mean that we can affect mobilization. For example, there remain important open questions about the generalizability of these findings. The results in this paper do not tell us, *prima facie*, if the particular innovative methods of organizing labor that Wikipedia used are likely to work for other types of creative goods or in projects with very different goals. Additionally, these results suggests limits in terms of what nascent projects set out to do. Many efforts that might seek to learn from these results to increase the likelihood of their own success will begin with goals in mind. Advice to “do something else,” or to “try to build a different type of product,” may be neither actionable nor useful.

These results may be both valid and limited in predictive power. In experiments around music popularity, Salganik et al. (2006) have shown that it is effectively impossible to predict the success of songs by looking only at qualities of the songs themselves because popularity is largely driven by social interactions that are not caused by qualities of the music. Predicting the success of firms or products seems similar, at least superficially, to the goals of this study and it is notoriously difficult to build strong predictions (e.g., Watts, 2011). Quantitative tests of these hypotheses might fully confirm that projects with a familiar product and innovative process are more likely to mobilize volunteers than products that are innovative on both dimensions – but the effect size, in absolute terms, might be small. Although the proposed relationships might hold, on average, projects that “do everything right” may still be overwhelmingly likely to fail for reasons outside of initiators’ power.

With these limitations foregrounded, this paper offers an important set of contributions to the literature on peer production. From a methodological perspective, this paper offers an answer to the critique that the peer production literature has spent too much time studying successful projects. If peer production scholars care about mobilization, this pa-

per reflects a step forward in that it stops selecting on the dependent variable and, through that process, brings relative failures into the analysis. More importantly, it offers a theory about why some attempts to build peer production projects lead to large mobilizations, like Wikipedia and Linux, while the vast majority struggle to attract even a second contributor. It offers a step, supported in empirical evidence, toward a broader causal theory of mobilization.

Essay 2

*The Remixing Dilemma: The Trade-off Between Generativity and Originality*¹

2.1 INTRODUCTION

Remixing, the reworking and recombination of existing creative artifacts, represents an enormous, important, and controversial form of online engagement. Most commonly used in reference to the creation of music, video, and interactive media, Manovich (2005) has called remixing, “a built-in feature of the digital networked media universe.” Lessig (2008) has argued that remixing reflects both a broad cultural shift spurred by the Internet and a source of enormous creative potential. Benkler (2006) has placed remixing at the core of “peer production” – the organizational form behind free and open source software and articles on Wikipedia – and has argued for the deep cultural importance of remixing. Scholars of innovation have suggested that remixing practice plays a critical, if under-appreciated, role in new forms of innovation facilitated by the Internet (e.g., von Hippel, 2005).

To advocates of remixing, and to proponents of peer production more generally (e.g., Raymond, 1999; von Krogh and von Hippel, 2006), the fecundity or “generativity” of creative works is of utmost importance in that it determines remixing’s very existence. But although collaboration lies at the heart of definitions of peer production, and despite

¹This essay is collaborative work with Andrés Monroy-Hernández. A version of this essay was published in the *American Behavioral Scientist* in May, 2013. As per the copyright agreement, this text is based on the accepted version of the article. The published version can be found at: <http://dx.doi.org/10.1177/0002764212469359>

the enormous amount of collaboration that occurs in some of its poster children, most articles on Wikipedia and other wikis never attract many editors (Ortega, 2009; Kittur and Kraut, 2010), most Free and Open Source Software (FLOSS) projects founder (Healy and Schussman, 2003), the majority of YouTube videos are never remixed, and most attempts at “meme spreading” on 4chan fall flat (Bernstein et al., 2011).

Proponents of remixing argue that generativity leads to increased innovation and democratized production. For example, Zittrain (2008) argues that some technologies, like the Internet, are generative and important not because they solve problems directly but because they provide rich and unconstrained platforms upon which derivative technologies can be built. Previous research has looked at ways to promote remixing, especially among young people (e.g. Cheliotis and Yew, 2009; Luther et al., 2010; Hill et al., 2010). Jenkins et al. (2006) have argued that educators can work to increase remixing behavior in young people.

Although remixes are defined by their derivative nature, the promise of remixing is contingent on the *originality* of derivative works. Remixing in the form of near-perfect copying seems unlikely to achieve Benkler’s (2006) goal of “making this culture our own,” or in building the transformative and empowering improvements at the heart of Zittrain’s examples. We also know users of remixing communities react negatively to visibly similar remixes of their projects (Hill et al., 2010). Moreover, issues of originality are often at the center of moral and legal discussions of remixing (Aufderheide and Jaszi, 2011). For example, to receive protection under US copyright law, a derivative work must be original in the sense that it is, “independently created by the author (as opposed to copied from other works), and that it possesses at least some minimal degree of creativity” (O’Connor, 1991). A critic of peer production, Keen (2007) conflates remixing on the Internet with copying saying, “the pasting, remixing, mashing, borrowing, copying – the stealing – of intellectual property has become the single most pervasive activity on the Internet.”

In this paper, we look at behavior within a peer production community to try to understand how designers of peer-production systems might, or might not, be able to support remixing that is both *generative* (i.e., likely to engender derivative works) and *original* (i.e.,

derivative works differ substantially from their antecedents). We attempt to answer two related research questions. First, what makes some creative works more generative than others? Second, what makes some creative works engender more transformative derivatives? Previous descriptive and theoretical work on peer production has pointed toward answers to both questions but has largely eschewed testable theories and hypotheses. Elaborating on foundational theory in peer production, and supported by empirical tests, we suggest that the answers to these two questions point in opposite directions and imply a trade-off for designers seeking to support remixing in online communities. *We suggest that three factors associated with higher levels of generativity – moderate complexity, creator prominence, and cumulativeness – are also associated with decreased originality in the resulting remixes.*

Using data from Scratch – a large online remixing community where young people build, share, and collaborate on interactive animations and video games – we present evidence that supports and extends several widely held theories about the foundations of generativity and originality. Our results suggest that designers of online collaborative communities may face a dilemma obscured by those celebrated exemplars of peer production communities: that system designs that encourage and support increased rates of remixing may also result in more superficial products.

In Section 2.2 of this paper, we discuss theoretical scholarship on remixing and motivate a series of six hypotheses about the predictors of generativity and originality, each of which we state explicitly following a description of relevant literature. In Section 2.3, we present our empirical setting and analytic strategy for testing these hypotheses. In Section 2.4 we present the results of our analysis and in Section 2.5 we discuss a series of important limitations of our findings and several tests that suggest that these limitations do not drive our results. We conclude in Section 2.6 with a discussion of future work and implications for the designers of interactive systems.

2.2 BACKGROUND

Generativity

Because most research on online peer production has focused on the most successful projects (see Crowston et al., 2010), we still know very little about why some peer production efforts become highly generative while the vast majority never attract contributors. Although foundational theories in peer production offers guidance, we must first elaborate on these theories to formulate testable hypotheses about the antecedents of generativity.

Zittrain (2008) posits the “Principle of Procrastination” that proposes that generative technologies tend to be designed in a way that leaves most details for later saying, “generative systems are built on the notion that they are never fully complete, that they have many uses yet to be conceived of, and that the public can be trusted to invent and share good uses.” For example, Zittrain suggests that the Internet was a more effective platform for innovation than corporate networks like *Prodigy* and *CompuServe* because its relative simplicity offered fewer constraints for potential innovators. In his influential essay, *The Cathedral and the Bazaar*, Raymond (1999) suggests that FLOSS projects like Linux attract participants because they “release early, release often” – that is, they publish their code earlier encouraging more collaboration in the form of feedback, bug fixes, and improvements.

Although there are several possible mechanisms through which “procrastination” and early releases might lead to generativity, one mechanism is the relative simplicity of these works. Early stage and incomplete projects will be simpler and easier for would-be contributors to understand and build off. Because these earlier, less complete, works are buggier, or more open-ended, they may also offer more avenues for engagement.

But while we interpret theory as suggesting that increased simplicity will be associated with generativity, this seems unlikely to hold for extremely simple works. The earliest possible release of Linux would, by definition, do nothing. It seems very unlikely that a featureless or extremely broken operating system kernel would excite and elicit contributions from other programmers in the way that Linux did. Similarly, if the designers of the Internet procrastinated on *all* features and created nothing, it seems very unlikely that

their system would have been an even more generative platform.

Hypothesis 1A: There will exist an inverse-U-shaped relationship between a work's complexity and its generativity.

Exposure to a work is, by definition, related to its generativity in that a work has to be seen to be remixed. Theorists have suggested that the relationship between popularity and remixing may run deeper. The antecedents of remixes, unlike some other forms of peer production, almost always have identifiable authors (Sinnreich, 2010). Lessig's (2008) key examples include music videos based on widely popular news footage and popular music and films. In Lessig's account, the act of remixing is often understood as a social statement of parody or critique. Jenkins (2008) documents how youth use fan fiction to create remixes of popular and culturally salient products and symbols.

Within particular communities, research has suggested that more popular individuals attract more remixers (Cheliotis and Yew, 2009). Using surveys and interviews with musicians, Sinnreich (2010) suggests that remixing is about creating explicit connections with previous, culturally salient, creators and that, "mash-ups are premised on the notion of recognizability and critique of pop culture," and that mash-up artists avoid rare vinyl samples in favor of popular songs to maintain these connections.

To the extent that remixing relies on cultural salience, we expect works of greater salience to be more generative. However, since popularity of the work itself might simply measure exposure, we should operationalize salience by looking to the prominence or "fame" of a work's creator while controlling for the exposure of the work in question. In other words, after having been viewed the same number of times, we expect a work by a more prominent creator to be more generative than a work by a less prominent author.

Hypothesis 1B: The prominence of a work's author will be positively related to its generativity.

We suggest that a third determinant of generativity is "cumulativeness:" a term we borrow from Murray and O'Mahony (2007) to describe works that aggregate the efforts of

many individuals through accretion and accumulation. Cheliotis and Yew (2009), Healy and Schussman (2003), and others have shown that activity in peer production is distributed unequally and that only a very small number of peer production projects incorporate the work of a large number of individuals building on one another's efforts. The majority of efforts are largely, or even entirely, uncollaborative.

Cheliotis and Yew (2009) have suggested that highly unequal rates of collaboration among projects in the ccMixer community is driven by a process of "preferential attachment" (Barabási and Albert, 1999) or cumulative advantage (DiPrete and Eirich, 2006) where, "works exhibiting a high degree of reuse become more attractive for further reuse." Cheliotis and Yew also suggest that, with important limits, remixing behavior will tend to form "chains" of remixed-remixes. To the extent that ccMixer is representative of other peer production projects in that collaboration drives more collaboration, we expect that cumulative remixes will be more generative than non-cumulative *de novo* works.

Hypothesis 1C: Works that are remixes themselves will be more generative than de novo projects.

Originality

Although theory on the relationships between remixed media, its creators, and the nature of their remixes is less developed, we find justification in existing theory for three hypotheses that parallel our hypotheses about generativity. In all three cases, we believe that theory points toward hypotheses suggesting that the qualities associated with higher generativity are also associated with lower originality in the resulting remixes.

In the previous section, we hypothesized that one possible mechanism behind "release early, release often" and the Principle of Procrastination is that simple projects are easier for new contributors to build on. Zittrain suggests that the generativity of a work will be determined, in part, by how easily new contributors can master it. Based on this, we posit that moderately simple works might be more generative than more complicated or simpler works because they are accessible to a relatively larger group of potential remixers. Driven by a marginal decrease in the effort or skill necessary to remix, we suggest that the remixes

produced through this process will also involve less effort or skill and, as a result of these linked processes, will be relatively less transformative.

Although it is also possible that more complex works are closer to “completion” than relatively simpler works and are therefore subject to less intensive improvements, we suggest that originality in remixes will be driven primarily by wider participation in the act of remixing, and that, as a result, remixes of works of intermediate complexity will tend to be less original than very simple or very complicated works.

Hypothesis 2A: There will exist a U-shaped relationship between a work’s complexity and the originality of its derivatives.

When discussing generativity above, we hypothesized that the creation of remixes of highly prominent creators is one way that remix artists seek cultural resonance for their works. To achieve this, it is important that a remixer maintain the recognizability of the original. For example, several musicians interviewed by Sinnreich suggest that P. Diddy’s song *I’ll be Missing You* became a cultural and commercial success in part because it consisted largely of a minimally modified version of the 1983 song *Every Breath You Take* by the band The Police.

Sinnreich argues that highly derivative remixes of culturally salient works strive to maintain a high degree of recognizability with the antecedent in the remix. To the extent that remixing of prominent work is more likely to be a form of cultural conversation, we will also expect the remixes of more popular or culturally salient works to be remixed lightly. On the other hand, when remixing the work of less prominent creators, the choice of a particular work might be driven more by use-value and, as a result, recognizability may play a less important role.

Hypothesis 2B: The prominence a work’s author will be negatively related to the originality of its derivatives.

Raymond (1999) describes “Linus’ law” – “with enough eyeballs, all bugs are shallow” – to suggest that collaboratively produced software will be higher quality and less buggy.

Benkler’s theory of peer production suggests that it is lightly motivated individuals contributing small amounts who participate in some of the most collaborative, and most cumulative, works of peer production. As a result, we might expect cumulative remixes (i.e., remixes of remixes) to begin with a less buggy or more complete work and, as a result, have less work to do.

Also suggesting a limit to generativity, Cheliotis and Yew (2009) observe that when a project is very cumulative and the product of many subsequent reuses, it becomes less likely to be reused in future generations. To the extent that the “chain” network structure becomes decreasingly likely to continue as it grows in length, we might expect that the existence of a shared goal (stated or implicit) for cumulative work may influence its continuation. Although Cheliotis and Yew do not present data on the originality of remixes, one explanation of their observation on chain remixes is that cumulative remixing will, on average, represent a process of refining and elaborating that has limits.

Hypothesis 2C: Remixes of works that are remixes themselves will be less original than remixes of de novo projects.

We do not suggest that these six hypotheses reflect a complete theory of generativity or originality in peer production communities. These hypotheses reflect our attempt at a partial theory in that they attempt to highlight three of the most widely cited theoretical determinants of generativity and originality. To our knowledge, none of these hypotheses has been tested empirically.

2.3 EMPIRICAL SETTING AND METHODS

Scratch

To test our hypotheses, we turn to the Scratch online community: a public and free website with a large community of users who create, share, and remix interactive media. The community is built around the Scratch programming environment: a freely downloadable desktop application, akin to Adobe Flash, that allows amateur creators to combine images, music, and sound with programming code (Resnick et al., 2009). A screenshot of

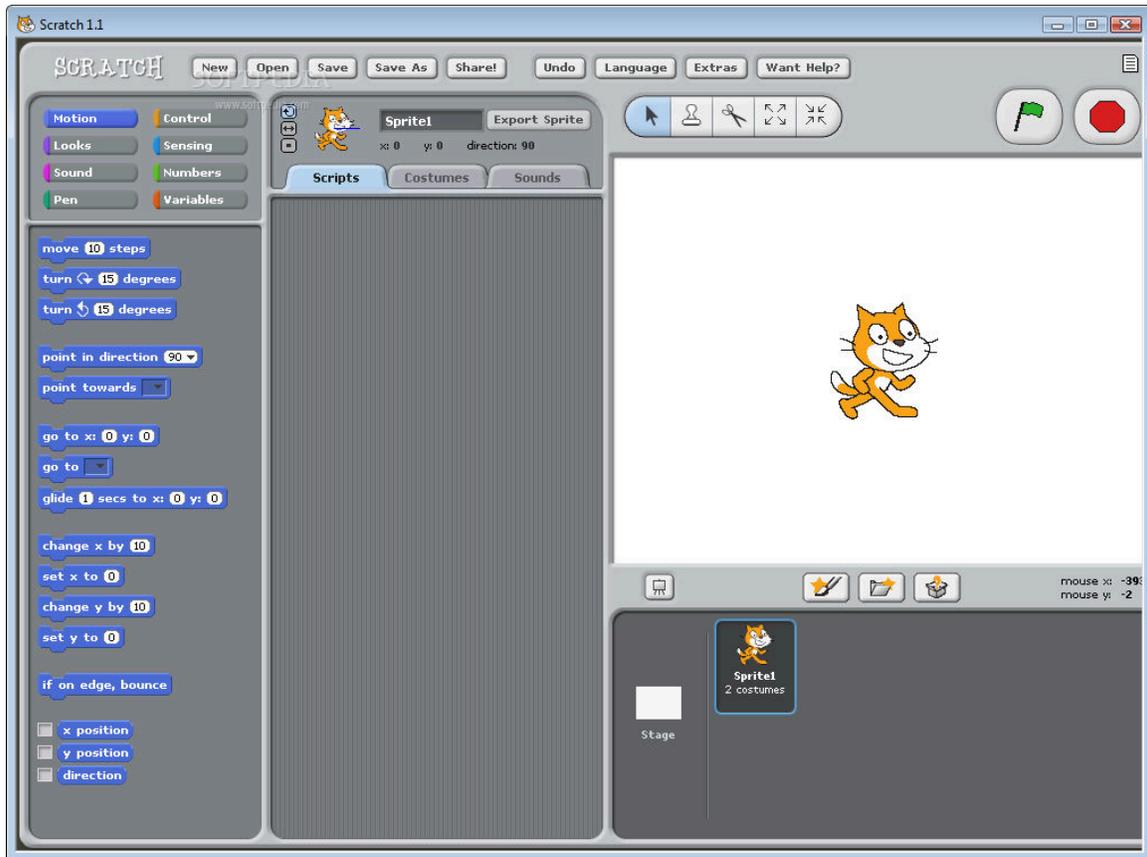


Figure 2.1: Screenshot of the Scratch desktop application. The leftmost column represents an inventory of possible programming blocks (shown in blue) which are assembled into program code in the center column. The white area in the top right represents the project as it will be displayed to a user interacting with the finished product. The bottom right column shows the available images (i.e., sprites) which are controlled by the code.

the Scratch programming environment is shown in Figure 2.1. Scratch was designed by the Lifelong Kindergarten Group at the MIT Media Lab as a platform for constructionist learning (e.g. Papert, 1980) and aims to introduce young people to computer programming. Scholars have located much of the practice and promise of remixing in communities of young technology users (e.g., Lessig, 2008; Jenkins, 2008; Palfrey and Gasser, 2008). With a large community of young users, Scratch represents an ideal platform to study remixing.

From within the Scratch authoring environment, creators can publish their projects on the Scratch community website hosted at MIT.² As of April 2012, more than one mil-

²<http://scratch.mit.edu>.



Figure 2.2: Screenshot of the Scratch online community front page as it appears to users visiting it from the Internet. The page is dynamically updated every several minutes to highlight new activity. (Accessed: April 1, 2012)

lion users had created accounts on the website and more than one third of these users had shared at least one of more than 2.3 million total projects. As the only web community built around sharing Scratch projects, the community contains virtually all Scratch projects shared online.

The nature of Scratch projects varies widely and includes everything from interactive greeting cards to fractal simulations to animations to video games. The community is

visited by more than half a million people each month³ who can browse material on the website but visitors must create accounts in order to download projects or contribute in the form of publishing, commenting, showing support, tagging, or flagging projects as inappropriate. A majority of the community's users self-report their ages ranging between 8 to 17 years old with 13 being the median age for new accounts. Thirty-five percent of users of the online community self-report as female.

Central to the purposes of this study, the Scratch online community is designed as a platform for remixing. Influenced by theories of constructionist learning in communities (Papert, 1980) and communities of practice (Lave and Wenger, 1991), the community seeks to help users learn through exposure to, and engagement with one another's work. The commitment to remixing is deep and visible in Scratch. The name "Scratch" is a reference to hip hop disc jockeys' practice of remixing. Every project shared on Scratch is available for download and remix by any other user, through a prominent download button. Additionally, every project is licensed under a Creative Commons license – explained in "kid friendly terms" – that explicitly allows reuse. Administrators and community members routinely encourage remixing.⁴

Issues of generativity and originality play a prominent role in the Scratch community. Although Scratch is designed as a remixing site, experience suggests that only around one tenth of all projects of Scratch projects are likely to be ever be remixed. Approximately 2% of remixes are flagged as inappropriate – often with accusations of unoriginality like, "This is MY own artwork he has uploaded without an ounce of originality." On the other hand, Scratch creators often explicitly encourage others to remix their works and even request help creating new features and solving bugs. Scratch users frequently respond to these requests with remixes but also frequently remix without prompting or communication.

Data and Measures

The Scratch online community is built on top of a large relational database that contains a wide variety of metadata on projects, users, and interactions on the website. Crucially,

³<http://quantcast.com/scratch.mit.edu>

⁴http://info.scratch.mit.edu/License_to_play

the Scratch website identifies, tracks, and presents data on whether and when projects are created via remixing. Additionally, the website stores each of the “raw” Scratch project files which can be further analyzed to reveal details, such as each project’s programming code. Our dataset is constructed both by exporting metadata about the community’s users, projects, and interactions, and by detailed algorithmic analyses of each project to compute the differences in code between remixes and their antecedents.

Our unit of analysis is each Scratch work or project (p). Our dataset consists of the 536,245 projects shared in the Scratch Community in a one-year period over 2010.⁵ We do not include data on 1,182 projects which we were not able to analyze due to technical errors in our tools or corruption in the project files. We also omit 136,968 projects (21%) that were removed from the site by their creators, although we find, in robustness checks not reported here, that our results are not substantially different when we include them. The projects in our dataset were shared by 105,317 unique users (u). We selected data from 2010 because Scratch’s administrators felt that the site and its community were mature and stable during this period and because the Scratch website did not undergo any significant design changes that might have affected remixing behavior. Because our data is longitudinal, we track each project for one year and, for time-varying measures, present measures at the end of the one year period.

To answer our first three hypotheses, we operationalize generativity in two related ways. First, we construct a count of the number of remixes of each project shared within the first year of the project’s publication ($remixes_p$). Second, we construct a dummy variable indicating if a project has been remixed by another user at least once in the one year period subsequent to being shared ($remixes_p > 0 \Rightarrow remixed_p = 1$). Because multiple-remixing is often endogenous (i.e., individuals may choose to remix a previous remixed project because others have remixed it), we suggest that our dichotomous measure ($remixed$) offers a more reliable, if more conservative, measure of generativity.

We operationalize complexity of projects as blocks ($blocks_p$). Blocks, shown in Figure

⁵Due to the timing of data collection, we use all the projects that were created during a one-year period from December 1, 2009 through December 1, 2010. We follow each project from this window for one year and our last data collected is from December 1, 2011.

Variable	N	Mean	SD	Min	Max
Dependent Variables					
Remixes > 0 times in 1 yr. (<i>remixed_p</i>)	536245	0.07	0.26	0	1
Remixes within 1 yr. (<i>remixes_p</i>)	536245	0.15	1.78	0	658
Edit Distance (Mean) (<i>distance_p</i>)	37512	85.57	397.66	0	21970
Question Predictors					
Number of blocks (<i>blocks_p</i>)	536245	99.60	476.19	0	196509
User’s cumulative views (<i>userviews_{up}</i>)	536245	1563.59	5546.90	0	197844
Remix status (<i>isremix_p</i>)	536245	0.18	0.38	0	1
Controls					
User age in years (<i>age_{up}</i>)	523092	17.57	11.62	4	74.75
Account age in months (<i>joined_{up}</i>)	536245	4.79	7.18	0	45.43
User is Female (<i>female_u</i>)	536222	0.37	0.48	0	1
Blocks per sprite (<i>blocks/sprites_p</i>)	536245	11.82	22.75	0	3111.50
Views within 1 yr. (<i>views_p</i>)	536245	13.57	69.90	0	4977

Table 2.1: Summary statistics for variables used in our analysis. Measures with the subscript p are measured at the level of the project while measures with the subscript u are measured at the level of the user.

2.1, are analogous to tokens or symbols in source code for computer programs. Blocks are similar to, but more granular than, source lines of code which have a long history of use as a measure of both complexity and effort in software engineering (Walston and Felix, 1977; Albrecht and Gaffney, 1983). Scratch projects may also contain media elements. Several possible metrics of media complexity (e.g., the number of images or sounds in a project) are highly correlated with (*blocks_p*). For that reason, because an integrated code and media measure is not available, and because we feel that a code-based metric is more granular and comparable across projects, we leave exploration of media-based complexity as an area for future work.

Data on user interactions in Scratch provide a series of possible measures of user prominence within the community. Possible indicators of prominence include the total past views of a user’s projects by other users, the number of past expressions of admiration or “loveits” (analogous to “liking” something on other social media platforms), total previous selections of a user’s work as another’s favorite, and total past downloads – each variable

is measured at the level of the project or work but is aggregated for each project’s creator at the point in time that a project is shared. Because these measures are highly correlated ($0.84 < \rho < 0.97$), we operationalize prominence as a user’s cumulative previous views ($userviews_{up}$) at the moment that the project in question was uploaded. Our results are similar using the other indicators. We operationalize cumulativeness using a dummy variable that indicates whether a project itself is a remix of another project ($isremix_p$).

Finally, we include a series of control variables that may also be associated with the generativity of projects and with the originality of subsequent remixes. For each user, we include self-reported measures of gender which we have coded as a dichotomous variable ($female_u$), date of birth which we have coded as age in years at the moment that each project was shared (age_{up}) which may indicate sophistication of the user, and age of each account ($joined_{up}$) which may indicate the level of experience of a user with Scratch. For each project, we are concerned with the effect of exposure on the likelihood of remixing, so we attempt to control for views using the number of times that each project was visited in its first year on the site ($views_p$).⁶ “Sprites” are the objects in Scratch project to which code is attached. Because more modular projects may be easier to remix, we also calculate a measure of the average numbers of blocks per sprite ($blocks/sprites_p$) which may act as a very coarse measure of modularity.

To answer Hypotheses 2A-C about originality, we create a new dataset that includes only the subset of 37,512 projects that were shared in the community during our one-year window and that were remixed at least once in the following year.⁷ We operationalize the originality of a remix using a calculation of the degree to which a project diverges from its antecedent. To calculate this divergence, we begin with the list of remix-antecedent pairs. Next, we identify and compare each code component of the remix to the corresponding code component in the antecedent. Our measure of originality is the Levenshtein “edit distance” (Levenshtein, 1966).

⁶As an alternate specification, we instead control for the number of unique users who have views the site, with nearly indistinguishable results.

⁷Due to technical errors or corrupted project files, we do not include 1,217 projects that site-metadata indicates were remixed but that we were unable to analyze. We believe that these errors were due to random corruption and are unlikely to bias our results.

Levenshtein distance is a metric that has been used widely in software engineering to measure the divergence of code. The traditional Levenshtein analysis is a character-by-character comparison. In our case, we use blocks as tokens and our measure of distance is the sum of distances across all code and represents the minimum number of changes to blocks that would be needed to convert an antecedent project into its remix. This metric come with several important caveats. First, the measure will not capture artistic charges. If every media element in a project were changed but the code left intact, our analysis would consider the projects perfect copies. That said, we draw some confidence from the fact that exploratory analyses suggest that media derivativeness and code derivativeness are highly correlated. Additionally, the measure does not reflect “conceptual” remixing – such as employing Disney or Nintendo characters in a new Scratch game. We hope to address these limitations in future work.

Of course, a given project can be remixed multiple times. In fact, 11,704 of the projects remixed in our window (31%) were remixed more than once within a year of being shared. The distribution of remixes was highly skewed: the maximum number of remixes in our sample was as high as 658, and the mean was 2.14. As a result, our measure of edit distance is the mean edit distance of all projects shared in the year following a project’s publication on the website ($distance_p$).

Analysis

Our analytic strategy involves the estimation of a series of two sets of parallel regression models. In both cases, we include variables operationalizing project complexity, creator prominence, and project cumulateness that correspond to our three sets of parallel hypotheses. Both *blocks* and *userviews* are highly skewed but a started log transformation results in an approximately normal distribution in each case. Hypothesis 1A and 2A predict a curvilinear relationship between the dependent variables and our measure of complexity. As a result, we include a quadratic specification for log *blocks* in each model and focus our interpretation on the coefficient associated with the quadratic term which will determine the direction of the curve. Because the amount of code in a remix does not reflect the

work of only the person sharing the project, we include an interaction term between our measure of *blocks* and *isremix* to capture the difference in the effect of complexity between remixes and *de novo* projects.

Providing tests of Hypotheses 1A-C about generativity, our first two models consider generativity in the full dataset of 523,069 projects shared in our window of data collection for which we have complete information.⁸ In our first and more conservative test, Model 1, we use logistic regression to model the likelihood of a project being remixed at least once on our sets of predictors and controls:

$$\begin{aligned} \text{logit}[P[\textit{remixed}_p]] = & \beta + \beta \log \textit{blocks}_p + \beta \log \textit{blocks}_p^2 + \beta \log \textit{userviews}_{up} + \\ & \beta \textit{isremix}_p + \beta \textit{age}_u + \beta \textit{joined}_{up} + \beta \textit{female}_u + \beta \log \textit{blocks}/\textit{sprites}_p + \\ & \beta \log \textit{views}_p + \beta (\log \textit{blocks}_p \times \textit{isremix}_p) + \beta (\log \textit{blocks}_p^2 \times \textit{isremix}_p) \end{aligned}$$

Model 2 also tests Hypotheses 1A-C using our second measure of generativity: the count of remixes of each project in the first year. It is otherwise identical to Model 1. Poisson regression is frequently used for count dependent variables but, as is common with counts, there is an over-dispersion of zeros in the number of times a project has been remixed. To address this overdispersion, we use a negative binomial regression strategy that estimates the right side of the equation in the model above on the count of *remixes*.

To test Hypotheses 2A-C about originality, we begin with a reduced dataset that consists of the subset of 36,722 projects which were remixed at least once after being shared, and for which we have the creator's age and gender data. The right side of Model 3 is, once again, identical to that of Model 1 shown above. The left side corresponds to the mean Levenshtein distance of every remix of the antecedent project. Because *distance* is a count and, like *remixes*, is overdispersed, we once again forgo Poisson regression in favor of a negative binomial count model.

	Generativity		Originality
	Model 1 (P[remixed])	Model 2 (remixes)	Model 3 (distance)
(Intercept)	-5.070*** (0.031)	-5.045*** (0.029)	2.437*** (0.053)
log <i>blocks</i>	0.525*** (0.016)	0.374*** (0.016)	-0.028 (0.027)
log <i>blocks</i> ²	-0.037*** (0.002)	-0.035*** (0.002)	0.052*** (0.003)
log <i>userviews</i> _{up}	0.023*** (0.003)	0.002 (0.003)	-0.041*** (0.005)
<i>is.remix</i>	0.786*** (0.045)	0.426*** (0.045)	-1.035*** (0.071)
<i>age</i>	0.000 (0.001)	0.007*** (0.001)	0.006*** (0.001)
<i>joined</i>	-0.006*** (0.001)	-0.006*** (0.001)	0.005*** (0.001)
<i>female</i>	-0.003 (0.012)	0.106*** (0.012)	-0.348*** (0.021)
log <i>blocks/sprites</i>	-0.517*** (0.012)	-0.375*** (0.012)	0.289*** (0.018)
log <i>views</i> _p	0.840*** (0.006)	1.028*** (0.006)	0.153*** (0.009)
log <i>blocks</i> × <i>isremix</i>	0.318*** (0.025)	0.303*** (0.026)	0.160*** (0.040)
log <i>blocks</i> ² × <i>isremix</i>	-0.045*** (0.003)	-0.032*** (0.003)	-0.002 (0.005)
θ		0.265*** (0.003)	0.301*** (0.002)
<i>N</i>	523069	523069	36722
AIC	219860.275	307810.178	313334.551
BIC	220396.313	308390.886	313777.130
log <i>L</i>	-109882.137	-153853.089	-156615.276

Standard errors in parentheses; * $p < .05$; ** $p < .01$; *** $p < .001$

Table 2.2: Model 1 is a logistic regression model of the likelihood of a project being remixed within one year. Model 2 is a negative binomial regression model of a count of the times a project will be remixed within a year. Both use the full dataset of projects ($N = 523,069$). Model 3 is a negative binomial regression model of a count of the mean edit distance for all projects remixed within a year of being shared ($N = 36,722$).

2.4 RESULTS

Model 1 seems to provide support for Hypothesis 1A; we see support for the inverted U-shape in the relationship between complexity and generativity in the negative coefficient on the quadratic term in Models 1 and 2 (see Table 3.2). Holding other variables at their sample mean, Model 1 predicts that a bit more than 1 percent of projects will be remixed at both the minimum (0 blocks), and maximum (196,509 blocks) in our sample. That said, this support comes with a critical qualification. For the vast majority of projects, marginal increases in complexity are associated with increased generativity.

An example can serve to illustrate this point. The distribution of projects by *blocks* is highly skewed toward more simple projects with the median project having only 26 blocks. In other words, although the most simple and the most complicated projects are indeed at lower risk of being remixed than projects of median complexity (i.e., in an U-shaped relationship), we estimate that projects have an increasing likelihood of being remixed into the 95th percentile of complexity. Even among very complicated projects, the relationship is effectively flat. For example, holding all other predictors at their sample means, Model 1 estimates that 6.77% of projects with 385 blocks (the 95th percentile) would be remixed while effectively the same proportion of otherwise identical projects with 1,204 blocks (the 99th percentile) would be.

As predicted in Hypothesis 2A, we see some support for an inverse-U-shaped relationship between complexity and originality by the negative sign on the quadratic log *blocks* term in Model 3. That said, there is little evidence that simple projects are associated with increased originality in remixes because the first-order term is not significantly different from zero, therefore suggesting a curvilinear relationship where edit distance is monotonically increasing with complexity over the range of our data. Indeed, holding other variables at the sample mean, Model 3 estimates that a project with 3 blocks (10th percentile) will have an average edit distance of 21 changed blocks while an otherwise similar project with 211 blocks (90th percentile) would, on average, be associated with an mean edit distance of 81 changes.

⁸We omit 13,176 projects for which we are missing age or gender data.

Tests of Hypotheses 1B and 2B on the relationship between author prominence and generativity and originality are given in the coefficients estimates associated with *userviews*: the log-transformed count of the number of times that other users have viewed a project's creator's work in the past. We find support for Hypothesis 1B on the positive relationship between author prominence and generativity. Holding other variables at their sample mean, Model 1 predicts that the odds of being remixed are slightly higher (1.02 times) for each log unit increase in the number of previous loveits the project's creator received, and that the result is statistically significant. Model 2 adds an important qualification to this support. There is no statistically significant association with prominence when operationalized as the number of remixes within one year. In other words, author prominence is a positive predictor of whether or not a project will be remixed but is not associated with higher numbers of total remixes.

We also find strong support for Hypothesis 2B that predicts a negative relationship between creator prominence and originality of remixes. Holding other predictors at their sample mean, we estimate that remixes of a project whose creator's previous projects had received no previous "views" (10th percentile) would have an average edit distance of 44 changed blocks. An otherwise identical project whose creator had received 3,652 previous "views" (90th percentile) would be estimated to have an average edit distance of 32 changed blocks.

Our results also provide strong support for Hypotheses 1C and 2C. Tests of the association between cumulativeness and measures of originality and generativity are captured by the parameter estimates associated with *isremix*. Model 1 suggests that the odds that a remix will be remixed is 2.2 times higher than the odds that an otherwise similar *de novo* project will be. Model 2 suggests that remixes will also be remixed more times. In strong support of Hypothesis 2C, Model 3 suggests that these remixed remixes will tend to be much less original. Holding other qualities at their sample means, Model 3 estimates that a remixed remix will have an edit distance of 24 changed blocks while a similar *de novo* project will have an edit distance of 41 changed blocks.

The models also include a statistically significant parameter estimate associated with

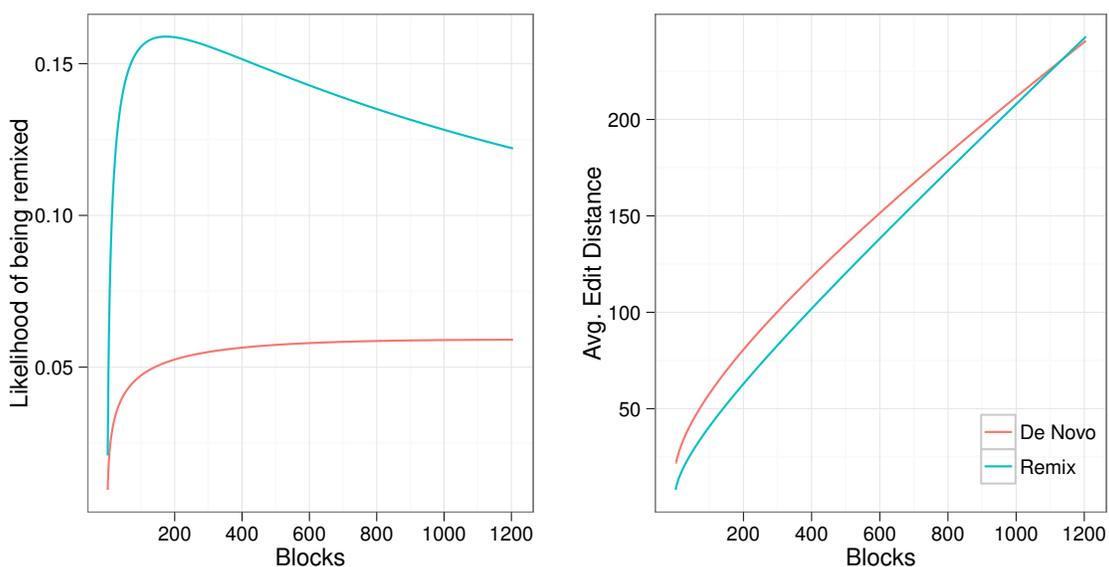


Figure 2.3: Two plots of estimated values for prototypical projects. Panel 1 (left) display predicted probabilities of being remixed as estimated by Model 1. Panel 2 (right) display predicted edit distances as estimated by from Model 3. Both models show predicted values for both remixes and *de novo* projects from 0 to 1,204 blocks (99th percentile).

the interaction between complexity and cumulativeness as measured in blocks. Because our measure of blocks is non-linear, interpretation of this result is complex. Prototypical plots of the estimates for remixes and non-remixes across our sample’s range of project complexity are shown in Figure 3.2. Both generativity models suggest that remixes are associated with higher rates of generativity for all projects in our dataset. Both models also suggest that the inverse-U predicted in Hypothesis 1A is likely stronger for remixes than *de novo* projects. Model 3 predicts that for nearly all projects in our sample, remixes of remixes will be less original, as measured by edit distance, than otherwise similar remixes of *de novo* projects, but that this difference is unlikely to be substantively meaningful.

2.5 LIMITATIONS

There are a number of important limitations and threats to the validity of the results presented above. First, although our data is longitudinal, our analytic strategy follows each project for one year and treats these data as cross-sectional. Model 1 treats projects as being

remixed only if they are remixed within one year of being published. Of course, projects can and are remixed after one year and there is a risk that our results are biased by the fact that these “late bloomers” systematically differ from other projects. We can model this threat using non-parametric Kaplan-Meier survival functions of projects’ likelihood of being remixed on a full dataset of Scratch projects. The resulting estimates suggest a rapidly flattening survival function. For example, 30 days after being shared, 8% of projects are remixed while after a full year, 10.8% are; only 2.4% more projects will have been remixed two years after that. As a result, we are confident that the analysis presented considers a large majority of remixing activity. Of course, this does not rule out the threat that these “late bloomer” projects are very unusual and our results may be biased through their omission. In other analyses not reported here, we use Cox proportional hazard survival and counting process models to estimate the instantaneous risk of projects being remixed using a random subset of 100,000 projects. The signs and relative magnitudes of the coefficients in our model are not substantively different from those presented here.

A second concern is that the use of average edit distance in Model 3 may lead us to conclude that more generative projects will tend to have lower edit distances simply because all projects are susceptible to unoriginal remixing and that being remixed more often puts projects at increased risk of these very simple remixes. We address this threat by re-estimating Model 3 using the highest edit distance of any remix as our measure of originality. The results of this model are largely unchanged from those reported in Table 3.2 and, indeed, are even stronger in the estimates of the effect of prominence and cumulativeness.

Third, there is an important concern with blocks as an indicator of complexity. As we have already suggested, blocks will not capture complexity in ways that do not involve programming, such as story-telling and visual arts. They can also be “cut and pasted” in a way that may not correspond to complexity through increased effort. It is possible that very complicated cut-and-paste projects are skewing the results for complexity. We can address this with a unique measure available in Scratch. Each Scratch project records the time and date every time that a user clicks the “save” button as well as the time that the user shares the project. We can use the time between the first “save” and the point at which

the user shared the project as a proxy for effort.

This alternative measure is noisy in the sense that some users may not share a project for hours, days, or weeks, but not spend that entire period engaged in work on the project. Additionally, 44% of the projects in our window were shared without ever being saved once, so values on this indicator are missing. With these limitations in mind, we re-estimate Models 1 and 2 on the subset of 298,926 projects for which we have data and replace our measure of blocks with “minutes-to-share” (*MTS*). We find that our results in modeling generativity using this alternative specification are essentially unchanged. A similar re-estimation of Model 3 using the 22,048 remixed projects with *MTS* data did not find support for either the quadratic specification of *MTS* or its interaction with *isremix*, but offered substantively similar predictions in its estimation of a positive linear association between edit distance and complexity, which leaves our findings essentially unchanged. This robustness check also give us additional confidence in the applicability of these results to media and other non-code complexity.

In other robustness checks, we add random effects to control for possible clustering due to the fact that a single user can upload multiple projects and find that our estimates and results are unaffected. We also use robust estimation of standard errors to address concerns of potential heteroscedasticity. Using robust estimates for Model 3, the interaction terms between *blocks* and *isremix*, already substantively similar, are rendered statistically insignificant. The rest of our results, and our findings for each of our hypotheses, are unchanged.

A final concern, common to studies of peer production, is the question of generalizability. We have tested generalizability to other periods of time within the life of the Scratch community and find that our results are similar. Although we cannot speak to the generalizability of these results to other remixing communities or peer production projects, we believe that remixing in Scratch provides insight into the behavior of young creators more generally. The degree to which these results will generalize to adults, to other communities, or to activities other than the creation of animations and games, remain largely open questions for future research.

2.6 DISCUSSION

This paper provides support for the following paradox: attributes of works associated with increased generativity are associated with decreased originality, and vice versa. Our findings are based on a set of six hypotheses built on foundational theories of peer production which are tested using data from the Scratch online community. We find at least some support for the hypotheses that a work’s generativity has a U-shaped relationship with its complexity (H1A), that it is positively related to the prominence of the work’s creator (H1B), and that it is positively related to the work’s cumulative nature (H1C). We also find at least some support for the hypotheses that the originality of remixes will have an inverse-U-shaped relationship to the complexity of the antecedent work (H2A), that it is negatively associated with the prominence of the antecedent’s creator (H2B), and that it is negatively related to the work’s cumulative nature (H2C).

This paper’s primary contribution for system design theory is the proposal of a critical trade-off between the quality and quantity of remixes. To the extent that these results generalize, designers may need to trade-off deeper remixing with increased collaboration. Our findings imply difficult decisions around manipulating the visibility of variables such as author prominence and project complexity. For example, designers of a new peer-production system in need of more content might want to build features that further emphasize the salience of author prominence and remix “chains” in order to encourage generative content. However, our findings suggest that these designs might come at a cost in terms of the originality of the derivative works. Our results regarding the relationship of complexity to generativity and originality of remixes suggest that supporting increased complexity, at least for most projects, may have fewer drawbacks.

Many social media sites, including YouTube and DeviantArt, track and display user prominence using a metric of aggregate views nearly identical to our operationalization of prominence. Other sites try to incentivize collaboration with prominence through leaderboards. Our results suggest this technique can lead to increased generativity but might also lead to a decrease in originality due to the incentive itself. We also suggest that it may be important to avoid rewarding correlates of generativity for their own sake

when it is generativity that a designer wants to encourage. Encouraging cumulateness by incentivizing or raising the visibility of cumulative projects may be another way for system designers to encourage generativity; but, once again, our results suggest it may also be at the expense of originality of the resulting remixing.

Surprisingly, our weakest support is for the hypotheses about complexity that stem from our elaboration of Zittrain’s “Principle of Procrastination” and Raymond’s exhortation to “release early and release often” – the most widely cited theories of generativity. We find support for our hypothesis that the most complex projects will be less generative than projects of moderate complexity, but only if we consider the very most complex examples. In general, we find largely positive relationships between complexity and both generativity and originality over most of our data. This may point in the direction of one potential solution to the “remixing dilemma” we propose. It may also be that the young users of Scratch are unlikely to create projects that are complex enough to trigger the effect suggested by theory. It may also be that complexity is simply a poor measure of completeness, earliness, or open-endedness as it is theorized by Zittrain and Raymond. More research is needed to clarify this relationship.

It also bears noting that designers seeking implications in our research are often working at the level of socio-technical systems while our investigation is focused on the qualities of content shared within a single system. By holding the system constant, we hope to offer a deeper understanding of social dynamics that is essential to the design of well-functioning systems. That said, our findings are no substitute for experimental validation with between-community or longitudinal experimental designs. We see these as rich areas for future research.

Of course, nothing we have shown devalues the promise of remixing in terms of peer production, culture, and innovation. Indeed, we believe that societies’ ability to harness the power of remixing is deeply important, but requires further analyses similar to ours. Though our results suggest that highly generative works that lead to highly original derivatives may be rare and difficult for system designers to support, we do not suggest that encouraging them is anything but a worthwhile, and critically important, goal.

Essay 3

Laboratories of Oligarchy?

*How The Iron Law Extends to Peer Production*¹

3.1 INTRODUCTION

Commons-based peer production – the distributed creation of freely accessible information resources like Wikipedia and free software – represents a model of collective action and common-pool resource 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. It has transformed the way that firms in many industries do business, shifted how politicians campaign for office, and changed the way that individuals share information (Benkler, 2006; von Krogh and von Hippel, 2006; Karpf, 2012; Kollock, 1999; Shirky, 2008).

A central puzzle motivating this study emerges from the egalitarian principles and organizational goals of many peer production projects. According to various accounts, peer production projects function as novel forms of participatory organization with a broad democratizing potential. They have inspired a wave of social movement activists and theorists to embrace networked technologies as paths toward participatory democratic organizations (Benkler, 2006; Castells, 1996; Fuster Morell, 2012; Hess and Ostrom, 2011; Tufekci and Wilson, 2012). However, despite these ideals, we know that many successful peer production projects also exhibit strong inequalities of both participation and gover-

¹This essay is collaborative work with Aaron Shaw.

nance (Healy and Schussman, 2003; Fuster Morell, 2010). It has been widely noted that contributions to peer production follow a “power law” distribution where a small group of participants make an enormous number of contributions (Healy and Schussman, 2003; Ortega, 2009; Viégas et al., 2007). Many peer production projects, like Linux and Ubuntu, explicitly reject the label “democratic” and describe their leaders as “benevolent dictators for life” (Ingo, 2006; Hill et al., 2008).

This conflicting evidence suggests that while peer production communities have inspired a wave of movements and theorists to embrace networked technologies as tools for creating participatory democratic organizations, the communities themselves may reproduce a pattern of behavior more consistent with Robert Michels’ “Iron Law of Oligarchy” (1915), which suggests 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.

Michels’ theory has spurred an extended debate in social movement research as examples and exceptions have been sought out and examined closely in domains including the US labor movement and church groups (e.g. Jenkins, 1977; Lipset et al., 1956; Voss and Sherman, 2000). More recently, peer production communities like Wikipedia have been described as robustly democratic and resistant to the iron law (Konieczny, 2009). These studies have generally tested, contradicted, or elaborated aspects of Michels’ claims in individual organizations or in small samples. In contrast, our research builds on other recent attempts to draw comparative inferences about organizational behavior through the analysis of a large sample of voluntary groups or movement organizations (e.g., Andrews et al., 2010; Fuster Morell, 2010).

We seek to understand whether, given the existence of both unequal participation and systems of hierarchical organization within many peer production projects, large peer production organizations tend towards decreasing levels of organizational democracy as they grow. Using exhaustive longitudinal data of internal processes drawn from 684 wikis that have grown large and complex, we adapt Michels’ iron law to the context of peer pro-

duction communities and construct a series of quantitative tests of the iron law. We define several measures of oligarchy and democracy by drawing on the rich tradition of research studying movement organizations, unions, and political parties originally inspired by Michels (1915). In contrast to previous ethnographic findings on Wikipedia from Konieczny (2009), we present quantitative evidence in support of Michel's iron law in peer production. As wikis get bigger and older, a small group of leaders, present at the beginning, tend to consolidate power as their interests diverge from those of the other members of the collective.

In the section that follows, we discuss the relationship between peer production, the iron law, and democracy and further describe the concepts and theories used in our analysis. We build up to a set of three hypotheses that suggest that peer productions will be subject to the iron law. Next, we introduce the setting for our research, wikis hosted by the hosting firm Wikia, and present a more detailed explanation of our measures and analytic methods. Subsequently, we present the results of several models that test our hypotheses. In the final section, we discuss our results and the reasons why they provide strong support for the idea that the iron law applies to peer production communities. We also consider the implications of these findings for theories of online collective action, participatory democracy, and the role of networked tools for collaboration in democratic movements.

3.2 BACKGROUND

Peer Production and Participatory Democracy

The rise of online collective action as a mode of economic production, political participation, and information sharing has inspired many scholars to approach networked collectives as novel forms of democratic organization with broad democratizing potential (Benkler, 2006; Castells, 1996; Fuster Morell, 2010, 2012; Hess and Ostrom, 2011; Kolllock, 1999; Weber, 2004). A growing body of 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; Polletta, 2013). Much of this research has assumed the democratizing effects of certain technologies. By the logic of

these assumptions, democratizing effects are attributed either to the intrinsic properties of new communication technologies, or to the modes of social interaction and organization afforded by these tools.

We seek to question these assumptions and to evaluate them empirically in a large sample of online collectives engaged in commons-based peer production (Benkler, 2006). There are many reasons why peer production communities are unlikely to provide a direct point of comparison for social movement organizations engaged in contentious politics – the focus of much work on technology and collective action to date. That said, to the extent that peer production projects represent a digital vanguard, they can help address the broader issue of the prospects for participatory democracy in digitally networked organizations as well as the implications of digital networks for participatory democratic movements.

In a recent review essay, Francesca Polletta (2013) catalogs some of the ways in which the Internet and networked collaboration among online groups have contributed to a flowering of novel concepts and practices of participatory democracy in the context of contentious politics. The organizations and movements embracing what Polletta describes as a new wave of “democratic enthusiasms” (p. 42) reflect the learning and experiences of the new left and subsequent movements that took up the standard of enacting deeply democratic social transformation through radically democratic processes. Within these collectives aspiring to a participatory democratic ideal, processes for managing consensus and distributing participation have evolved to address concerns raised by earlier critics like Freeman (1973) and Breines (1989). Numerous scholarly and popular accounts have argued that digital information and communication technologies have a transformative democratizing capacity attributed, in large part, to their ability to reduce the organizational and economic costs of coordination and communication (e.g. Bennett and Segerberg, 2012; Earl and Kimport, 2011; Bimber et al., 2012; Karpf, 2012; Polletta, 2013; Shirky, 2008; Tufekci and Wilson, 2012).

Peer production communities such as those engaged in the creation of Wikipedia and the GNU/Linux operating system have not only inspired a wave of scholarly optimism

about the opportunities for digitally-enabled participatory democracy, but have also provided inspiration for contemporary social and political movement organizations. From the WTO protests in 1999 to the Arab Spring and Occupy movements, the idea of adopting an “open source” approach to organizing and mobilizing has become a rhetorical aspiration of many recent movements (see Fuster Morell, 2012). One of the reasons for this aspiration stems from the apparent lack of hierarchical and bureaucratic control in Wikipedia and GNU/Linux. A popular narrative suggests that, in the manner of a digital vanguard, 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., Shirky, 2008).

Little previous empirical research on peer production has focused on explaining organizational governance, but the body of existing work points to this topic as a critical topic for comparative analysis (Fuster Morell, 2010; O’Mahony and Ferraro, 2007; Shaw, 2012). In order to focus on organizational-level variation across peer production communities, our research draws on approaches originally developed for studying firms, cooperatives, political parties, unions, and social movements. However, in the typology of organizations, peer production projects are unusual. They share characteristics with all of the organizational types we have mentioned as well as 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. Peer production projects and communities do not look or function exactly like other types of organizations. At the same time, they are organizations in the fundamental sense that they are collectivities with goals and boundaries – one reason we feel that organizational theory ought to inform our thinking about peer production and vice versa. As voluntary organizations, peer production projects like Wikipedia and GNU/Linux have had unprecedented success at mobilizing and retaining contributions to collective public goods. For this reason, we join social movement scholars and proponents of peer production in the belief that theories of voluntary associations and movement organizations are a salient reference point for research in organizational democracy.

The Iron Law of Oligarchy

To consider the prospects of participatory democracy among peer production projects, we turn to one of the most influential general theories of governance in voluntary organizations developed by the German sociologist and political theorist Robert Michels. Michels developed the theory of the “iron law of oligarchy” as a general explanation of why democratically-run political parties in Italy around the turn of the twentieth century became less democratic as they grew in size.² Michels argues that (1) organizational growth drives the transformation of an organization’s structure into an increasingly formal and complex bureaucracy giving rise to a small group of professional leaders who become increasingly isolated from, and independent of, their organization’s membership; and that (2) these increasingly isolated and less accountable organizational leaders develop independent interests in the preservation of the bureaucracy itself resulting in the transformation of the goals and activities of the organization. Following previous work, we refer to these distinct dynamics as the “structural” and “goal transformation” components of the iron law (Leach, 2005; Voss and Sherman, 2000).³

In order to test the iron law, we must first formalize and operationalize the two distinct components of the theory: the rise of oligarchic organizational leadership and conservative organizational goal transformation. To constitute an oligarchy, the governance and leadership of an organization should reside in the hands of a stable, entrenched, minority that exercises dominant – sometimes illegitimate – control over organizational resources and policy (Lipset et al., 1956).⁴ Conservative goal transformation concerns the impact of oligarchic organizational structure on the character of organizational policies and activities over time. Note that “conservatism,” in this formulation, may describe either a particular set of political and ideological positions or a shift towards organizational self-preservation.

²We do not attempt, in this essay, to give a comprehensive overview of research into the iron law and its application. In addition to Michel’s original text, we would refer readers to work by Jenkins (1977), Leach (2005), and Voss and Sherman (2000).

³It 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). That said, we incorporate several aspects of Leach’s argument, in particular her specification of oligarchy along the dimensions of legitimacy and formality.

⁴Compare to Leach (2005, pp. 316 and 329).

Often, studies of political parties, unions, and other politicized social movement organizations have collapsed these two notions of conservatism. Doing so makes sense in the context of organizations with transformative social and political objectives, such as those that inspired Michels. However, in the context of organizations that lack clear political objectives – such as commercial firms, church groups, or non-commercial collectives for production – it does not make sense to apply overtly political standards of conservatism. As a result, we follow other scholars in interpreting conservative goal transformation as describing a shift in organizational policies or activities that reflect an agenda of organizational self-preservation and oligarchic entrenchment.

Logically, the iron law may also imply a decline of indicators of organizational democracy. As a result, we draw on the definitions of organizational democracy developed first in the work of Lipset et al. (1956) who extended Michels' work in their in-depth study of International Typographical Union (ITU). Lipset et al. identified multiple factors that they argued contributed to the capacity of the ITU to “overcome” the iron law. While some of these factors are not applicable to our empirical setting, Lipset et al. emphasized the capacity of members with no leadership role to become involved in day-to-day activities within the organizations and to participate in political discussions and debates in the union.

Peer Production in Wikis and the Iron Law

With particular attention to Piven and Cloward (1977), Rucht (1999), and Voss and Sherman (2000), we build on previous work as we attempt to characterize the overall patterns of organizational democracy and oligarchy in a very large population of voluntary organizations. In particular, we focus on 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 over a network. 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. Wikis encompass many characteristics that make them comparable to

voluntary movement 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.

From the point of view of Michels' theory and previous work comparing voluntary and membership organizations, one of the most interesting qualities of peer production communities in general, and wikis in particular, is that their organizational form and governance have 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 real cost to additional leaders. 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, on the whole, intrinsically supportive of participatory democratic organizational forms, there is reason to believe that wikis will be among peer production's most democratic organizations.

At the same time, and despite the egalitarian ideals associated with peer production, 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 appear to 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) 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). 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. Unlike political parties or labor unions, wikis included in our sample were not, in general, formed with the objective of facilitating political or economic representation of their members. Most wikis are distributed groups of individuals brought together through the desire to exchange and share information about a topic.

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). Effective leaders within Wikis, and other peer production communities, embrace multiple styles and practices of authority (Zhu et al., 2012; Schweik and English, 2012). 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. These administrators occupy structural positions of authority inside wikis and, as a result, provide 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:

H1: *The probability of adding new administrators declines as a wiki grows.*

Second, we consider whether wikis experience conservative 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. As one measure of this concept, we examine a subset of “reverts” (i.e., complete removals of contributions) performed by the administrators of a wiki to other editors. While many reverts are appropriate rejections of low quality contributions or spam, they nevertheless reflect a means by which ad-

ministrators impose their vision of legitimate and valuable participation on others.⁵ More importantly, reverts provide a clear indication of the degree of divergence between the interests. 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 content on the project should be. Formally, we predict that:

H2: The number of reverts by administrators of contributions made by experienced contributors will increase as a wiki grows.

Finally, we also consider whether indicators of participatory democracy rise or fall as wikis grow. To do so, we translate the concepts of democracy developed by Lipset et al. (1956) involving participation in the associational life of an organization into the setting of peer production projects. To do so, we examine a standardized measure of the equality of contributions across editors in the form of a Gini coefficient of contributions to each wiki. Despite the stability of power laws of contribution, previous research has found that Wikipedia has become *more* equal in this measure over time (Kittur et al., 2007). Although the iron law offers few concrete predictions about the relative equality of participation in democratic organizations, it seems consistent with Michels' theory to suspect that participation would become less democratic. Formally, we predict that:

H3: The distribution of contributions across contributors will grow more unequal as a wiki grows.

3.3 EMPIRICAL SETTING

Our empirical settings is a population of wikis hosted by the for-profit firm *Wikia*. Wikia was founded in 2004, was directly inspired by Wikipedia, and sought to apply the Wikipedia model beyond the education-based scope of the Wikimedia Foundation (the organization that supports Wikipedia). Wikia was founded by Jimmy Wales, Wikipedia's founder,

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.

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 Social-Text) Wikia is unique in that it only hosts volunteer-produced, peer production projects. Although many organizations, teams, and classrooms use wikis to support cooperative work, Wikia will not host these "closed" communities and never restricts viewership or editorship of its websites. Wikia's revenue stems from advertisements that it displays to non-contributing visitors to its hosted wikis.

Wikia does not pay for content and relies on peer production to create 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, ownership of the copyright on wiki content remains with the contributors, but all material is licensed freely to the public. All Wikia content is released under the Creative Commons Attribution-ShareAlike license (the same license that Wikipedia uses) and is made available for download.

In terms of the content and scope, Wikia wikis vary enormously, addressing popular culture and fan culture topics as well as things like information sharing about subcultures, software, food, fashion, and more. Some of the largest wikis in our study 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. Most of the largest and most highly edited Wikia wikis revolve around "fan culture" (see Lewis, 1992; Ito et al., 2012). Many wikis are topic-specific encyclopedias and are closely modeled after Wikipedia but are not constrained by Wikipedia's "notability" guidelines that restrict which topics can be covered. A large proportion of Wikia wikis include the term "pedia" somewhere in their name. In analytic terms, these wikis constitute a population of peer production communities.

Variable	Minimum	Median	Mean	Maximum	SD
Edits	644	13473	53366	2303248	161542
Articles	183	3168	11167	1270640	52988
Editors	69	218	790	68222	3456
Reverted Edits	0	285	1442	122950	5882
Administrators	0	7	11	247	18
Age (Months)	6	46	50	74	11

Table 3.1: Summary statistics for all of the wikis included in our analysis. ($n = 684$)

3.4 DATA AND MEASURES

Our full dataset includes data on 76,473 wikis created before April, 2010. It consists of more than one terabyte of data and includes the full history of every revision, from both registered and non-registered users, for every wiki hosted by Wikia from the time of the 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 until 2010 when they stopped being produced due to technical limitations caused by the dataset’s growing size. 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 which is developed, in collaboration, with the Wikimedia Foundation. Identical software makes it possible to derive a set of comparative metrics for analyzing organizational governance and activity across the wikis in the population. As with FLOSS projects, there is a highly skewed distribution in activity among wikis and the vast majority of wikis are small and relatively inactive. Considering sample median values, an average wiki in our full dataset would contain 225 contributions to 217 pages by 6 unique contributors at the point of data collection in 2010.

Because it seems orthogonal to Michels’ theory to evaluate whether a very small wiki is oligarchic, this analysis is limited to a subset of large wiki communities. 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 include corrupted or invalid data (e.g., edits marked as occurring before the wikis were founded) leaving us with a total of 684

wikis in our final dataset. Summary statistics for this subset of wikis are reported in Table 3.1.

Our dataset is longitudinal and every contribution made to each wiki is recorded in our raw dataset with timestamps accurate to within a second. 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 147,157 such observations.

When considering the activity of particular individuals within wikis, a critical distinction is administrator status, introduced in our background section above, which we use as an indication of structural and formal authority. To extract these data, we wrote custom software to visit each wiki in our dataset and to create detailed data on when individuals were promoted to, or demoted from, administrative roles. 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 users in another week may not.

Since the 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.

To test our three hypotheses, we construct three measures to act as dependent variables. For our first hypothesis (H1) that the probability of a community adding a new administrator will decline as wikis 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 2,934 wiki weeks (2% of the wiki-weeks in our sample).

To test H2 that the number of reverts by administrators of edits made by registered

editors will increase as wikis grow, we must first identify edits that qualify as reverts. Reverts are 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 editors who are registered members of the community provide a rough measure of the extent to which administrator interests diverges from those members' interests. To capture this, we construct a variable $admin-reverts_{week}$ that is a count of the number of edits that were reverted by administrators in a week on a wiki. Because we do not want these measures to reflect spam and vandalism, we limit this count to those reverts of edits made by registered users of the wiki.

H3 provides a test of democratic participation and suggests that the distribution of edits will grow more skewed as a wiki grows. To test H3, we compute Gini coefficients for the distribution of edits across editors during each week. Gini coefficients are a widely used measure of concentration or inequality (Gini, 1997). A Gini coefficient of zero would reflect a situation in which every editor of the wiki edited the same amount; a Gini coefficient of 1 would reflect maximum inequality such as a single editor making all contributions during a week.

We also include a set of control variables in each of our models. To capture differences in activity within a week, 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

In order 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 or democratic as they grow. Because our dataset includes repeated observations for each wiki, we

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 users edits can be undone in ways that do not return to the page to its previous state, this reflects a conservative measure of reversion.

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. Each of our models is fit with a different dependent variable. In Model 1 (M1), we use a multilevel logistic regression strategy to estimate the probability of a new administrator being added. Models 2 and 3 are both hierarchical linear models on three different dependent variables: (M2) the log-transformed number of reverts of registered users by admins; and (M3) the Gini coefficient capturing the concentration of edits across editors.

The right side of each equation is structured around a base model in which every variable is measured at the level of the wiki week. The model includes our independent variables, a base set of controls, and our compound error term:

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

There is an additional concern with M3 which regresses our covariates on Gini coefficients. Because editors and editors are the parameters used to compute the Gini coefficients, they are strongly correlated with the outcome in M3 and we remove them from the fitted models. Although the magnitude of coefficients are changed by their inclusion, the patterns of results for M3 are unchanged when they are present.

As is typical in datasets from online communities, many of the variables capturing activities in wikis are highly skewed. In particular, our measures of accounts, editors, pages, administrators, edits and administrator-reverts are each highly skewed. A started log transformation results in an approximately normal distribution in each case. The relationship between time and our dependent variables is curvilinear. As a result, we include a quadratic specification of *week* in our models.

3.5 RESULTS

We find that the wikis in our sample provide little evidence of robust resistance to oligarchy. As predicted in H1, we find that among the wikis in our sample, leadership roles

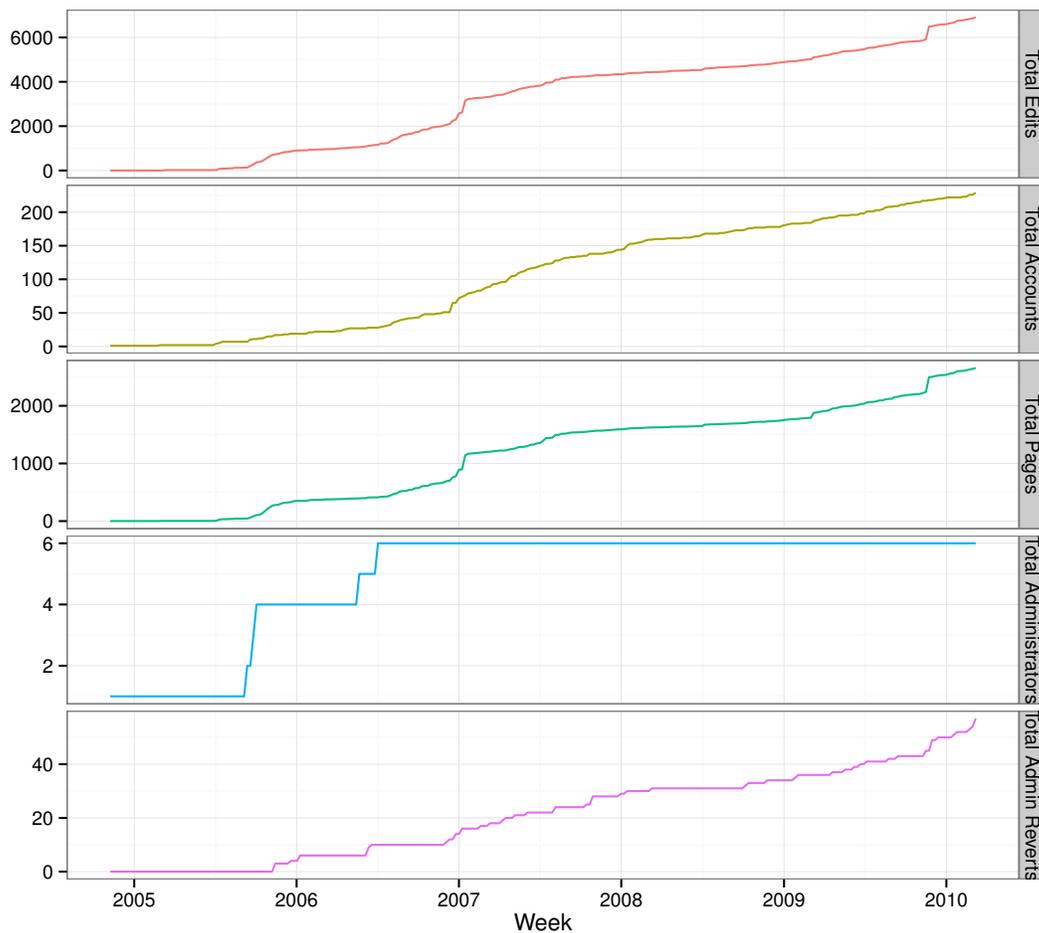


Figure 3.1: Growth plots for *Brickipedia*, an encyclopedia dedicated to LEGO products and designs and one of the online communities from our dataset.

and activities remained concentrated among an extremely small group of elites as the organizational membership grows. As predicted in H2, we found that as membership grows, the number of reverts of other users by the administration tended to rise with membership. As predicted in H3, we found that edits tended to become more concentrated and less equally distributed among editors. In other words, the wikis in our sample conform to a pattern predicted by Michel’s iron law.

Figure 3.1 provides an example of a single wiki from our dataset that reflects the basic pattern reflected in our models. The data shown are from *Brickipedia*, an encyclopedia devoted to products and designs from the Danish toy company LEGO. The top three panels

of Figure 3.1 show that the community has grown in terms of the number of contributions (*edits*), the number of registered users that have made at least one contribution (*accounts*), and the number of pages on the wiki (*pages*). Each of these variables 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 fourth panel that the number of *admins* has grown much more slowly and that no administrators at all have been added at all since mid-2006. The bottom panel shows that this small group of administrators has, over time, continued to remove contributions (i.e., “revert”) from other registered users.

Fitted regression models suggest a substantively similar pattern of associations and provide a set of formal tests for our hypotheses. Results from these models are shown in Table 3.2. In our first model (M1), we find strong support for Hypothesis 1 that larger communities are less likely to add new administrators. Indeed, we find that, *ceteris paribus*, one log-unit increase in the number of registered contributing users is associated with odds of adding a new administrator that 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 corp 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 strong support for Hypothesis 2 that as communities grow, the number of edits reverted by admins will also grow. Because both our dependent and independent variables are expressed in natural log transformed units, our parameters can be interpreted as elasticities. Again, we have included a series of important control variables. Most critically, we control for the number of edits in each wiki-week. At the margin, we estimate that a 1% change in the total number of accounts on a wiki, controlling for the total amount of activity in the wiki that week, is associated with a 5% increase in the the number of contributions made by registered users being reverted by administrators. In an alternative model, we remove our control for edits per week and predict estimated differences in the *proportion* of reverted edits by registered users. Our results

	M1	M2	M3
(Intercept)	−4.111*** (0.081)	−0.043** (0.015)	−0.182*** (0.004)
<i>week</i>	−0.006*** (0.001)	−0.003*** (0.000)	0.000*** (0.000)
<i>week</i> ²	0.000*** (0.000)	0.000*** (0.000)	0.000*** (0.000)
$\ln editors_{week}$	0.115** (0.038)	0.084*** (0.003)	
$\ln accounts_{total}$	−0.205*** (0.031)	0.050*** (0.002)	0.003*** (0.001)
$\ln pages_{total}$	−0.766*** (0.023)	−0.023*** (0.002)	0.003*** (0.001)
$\ln admins_{total}$	0.668*** (0.035)	−0.005 (0.006)	0.036*** (0.002)
$\ln edits_{week}$	0.994*** (0.022)	0.095*** (0.002)	0.130*** (0.000)
AIC	22163.868	222451.745	169197.999
BIC	22252.960	222550.736	169111.266
Log Likelihood	−11072.934	−111215.873	−84607.999
Deviance	22145.868	222431.745	169215.999
Num. obs.	147132	147132	113210
Num. groups: (<i>wikis</i>)	684	684	684
Variance: (Intercept) <i>wiki</i>	0.325	0.090	0.004
Variance: Residual		0.260	0.013

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

Table 3.2: Table of fitted multilevel regression models. The unit of analysis in each case is the wiki week. *M1* is a logistic regression regression model on the probability that a wiki will add a new administrator during a week. *M2* is a linear model of the number of reverts of registered users by admins over time. *M3* is a linear model of the Gini coefficient of edits among editors during the week-long period. ($n = 147, 157$ wiki weeks from 684 wikis.)

suggest a similar effect in either specification.

Our third and final model (M3) suggests that participation in wikis becomes less egalitarian as communities grow in complexity. Holding other variables constant, we estimate that wikis with one log unit more registered accounts will have Gini coefficients that are 0.3% higher. Although this effect is small, it is statistically significant.

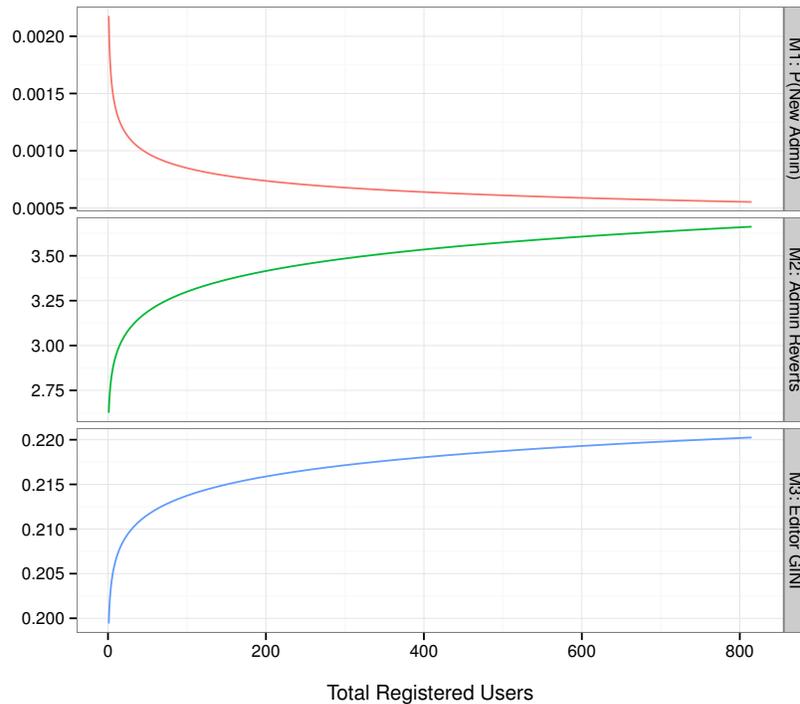


Figure 3.2: Plots showing predicted values from our models for wikis of varying size and complexity holding all other variables at sample medians. All outcome variables are measured in “per week” units.

To assist in further interpretation and comparison, plots of predicted values for prototypical wikis and for each of our three models are shown in Figure 3.2. 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 likely values of total registered accounts from 0 to 815 (the 95th percentile of observations in our dataset). These models show the predicted values for wikis that differ in terms of the number of total 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.

The plots emphasize what happens, on average, in our sample. One important limitation is that each of these plots controls for time. Although they do so at different rates, as organizations grow larger and more complex, they also grow older. Although prototypi-

cal plots 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 of a wiki constant.

The overall pattern of results provides strong evidence that, on average, as wikis become larger and more complex, a small group – present at the beginning – will restrict entry into positions of formal authority in the community while their interests diverge from the group and participation becomes more concentrated and unequal.

3.6 DISCUSSION

As we described in our introduction, recent theoretical and qualitative work on both collective action and peer production might lead one to expect that peer production communities will resist oligarchy and promote robust, egalitarian participation. Using a large sample of wikis, we present evidence to the contrary. The peer production communities we examine tend to reproduce non-democratic, non-inclusive organizational hierarchies of leadership and participation. As wikis grow larger, the probability of adding new leaders drops, entrenched leaders tend to use their authority to remove contributions of other community members, and the relative distribution of contributions within communities tends to become more unequal. The wikis in our sample are not indicative of robustly democratic and participatory institutions. This is true despite the relative lack of formal bureaucratic structure or clearly-defined roles. These results are consistent with Michels' iron law of oligarchy.

Our analysis remains limited and incomplete in many ways. In future work, we plan to incorporate additional controls as well as additional measures of oligarchy. We also plan to operationalize democracy more effectively with measures of the distribution of participation in deliberation and policy-making activity. 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 population of peer production projects. However, because of the size and scope of our dataset, we have confidence

that Michels' theory of organizational politics applies to peer production communities in several important respects.

Our results are particularly contingent for our third hypothesis that participation will become more unequal as communities grow. Research on remixing communities has suggested that as peer production projects increase in "cumulativeness," they will tend to attract less committed users who make more shallow contributions (Hill and Monroy-Hernández, 2013). In wikis, this may result in more inequality in participation. It is possible to imagine, for these and for other reasons, that editing on wikis may become more unequally distributed for reasons unrelated to oligarchy. More importantly, although we find an association between community size and inequality, our estimated effect is small in substantive terms. We present this result in context of this study because we believe that the finding points in the direction of a consistent pattern.

For scholars of politics, labor, and social movements accustomed to studying environments where the objective of mobilizing for collective action are often things like 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 influence that peer production technologies and organizational strategies have had on contemporary political movements. We also believe that even the most playful peer production communities face many of the same obstacles to collective action and public goods creation as social movements and other kinds of volunteer-based membership organizations.

This paper offers a series of contributions to existing research on participatory organizations. First, it joins a small number of studies in expanding the existing domain of organizational research beyond political parties, non-profits, and social movements in order to test some of the most influential and robust 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 scholarship that has historically not been applied to digitally networked groups. Third and finally, by conducting a comparative analysis across peer production communities in order to better understand

their social and organizational dynamics, we have also contributed one of the first tests of the widespread assumption that peer production projects inherently advance “small-d” democracy and are internally egalitarian.

Of course, some wikis in our dataset are more robustly democratic than others. Although the existence of these cases do not drive our findings, 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 an promising area for future research. But this research is precluded when scholars select on more democratic organizations or simply take the participatory and democratic nature of peer production organizations for granted.

We cannot know what these findings mean for the future of digitally networked social movements or democracy. However, we believe that our results suggest a provocative analogy with the sort of participatory associational culture many theorists of democracy speak of when they describe the foundations of more inclusive and just societies. For example, when Alexis de Tocqueville visited America in the early 19th century, he saw a flowering of civil society groups, the organizational structures of which he believed contributed to an inclusive, democratic culture (de Tocqueville, 2004). Similarly, when Seymour Martin Lipset and his colleagues looked to explain the remarkable democratic institutions within the ITU 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 have also brought with them the nationalization of labor politics as the organization of unions becomes part and parcel of the organization of national political culture.

The impact of online organizational platforms and communities in contemporary politics is unlikely to go away and some groups – like the Occupy movement – have explicitly

sought to model their movements on free software projects and other peer production communities. What sort of political culture and institutions will contemporary online associations, movements, and organizations create? Although, invoking de Tocqueville, peer production collectives have been treated as contemporary “laboratories of democracy,” our findings suggest that they may not necessarily facilitate enhanced practices of democratic engagement and organization. Indeed, our results suggest that widespread efforts to appropriate online organizational tactics from peer production may facilitate the creation of an entrenched form of oligarchy 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.

References

- Albrecht, A. J. and J. E. Gaffney (1983). Software function, source lines of code, and development effort prediction: A software science validation. *IEEE Transactions on Software Engineering SE-9*(6), 639–648.
- Andrews, K. T., M. Ganz, M. Baggetta, H. Han, and C. Lim (2010). Leadership, membership, and voice: Civic associations that work. *American Journal of Sociology* 115(4), 1191–1242.
- Aufderheide, P. and P. Jaszi (2011, August). *Reclaiming Fair Use: How to Put Balance Back in Copyright*. University Of Chicago Press.
- Barabási, A.-L. and R. Albert (1999, October). Emergence of scaling in random networks. *Science* 286(5439), 509–512.
- Benford, R. D. and D. A. Snow (2000, August). Framing processes and social movements: An overview and assessment. *Annual Review of Sociology* 26(1), 611–639.
- 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.
- Benkler, Y. (2013). Peer production and cooperation. In J. M. Bauer and M. Latzer (Eds.), *Handbook on the Economics of the Internet*. Edward Elgar.
- Bennett, W. L. and A. Segerberg (2012). The logic of connective action. *Information, Communication & Society* 15(5), 739–768.
- Bernstein, M., A. Monroy-Hernández, D. Harry, P. André, K. Panovich, and G. Vargas (2011). 4chan and /b/: An analysis of anonymity and ephemerality in a large online community. In *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media*, Barcelona, Spain, pp. 50–57. AAAI Press.
- Berzroukov, N. (1998, March). A second look at the cathedral and the bazaar. *First Monday* 4(12).

- Bimber, B. A., A. J. Flanagin, and C. Stohl (2012). *Collective action in organizations: interaction and engagement in an era of technological change*. Communication, society and politics. New York: Cambridge University Press.
- Breines, W. (1989). *Community and Organization in the New Left, 1962-1968: The Great Refusal*. New Brunswick, NJ: Rutgers University Press.
- Brown, G., T. B. Lawrence, and S. L. Robinson (2005, July). Territoriality in organizations. *The Academy of Management Review* 30(3), 577–594.
- Bush, V. (1945, June). As we may think. *Atlantic Monthly* 176(1), 101–108.
- Butler, B., E. Joyce, and J. Pike (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*, New York, NY, USA, pp. 1101–1110. ACM.
- Castells, M. (1996). *Rise of The Network Society (Information Age Series)* (1 ed.). Wiley-Blackwell.
- Charmaz, K. (2006, January). *Constructing Grounded Theory: A Practical Guide through Qualitative Analysis* (1 ed.). Sage Publications Ltd.
- Cheliotis, G. and J. Yew (2009). An analysis of the social structure of remix culture. In *Proceedings of the fourth international conference on Communities and technologies*, University Park, PA, USA, pp. 165–174. ACM.
- Coase, R. H. (1937, November). The nature of the firm. *Economica* 4(16), 386–405.
- Coombe, R. J. (1998). *The Cultural Life of Intellectual Properties: Authorship, Appropriation, and the Law*. Duke University Press Books.
- Corbin, J. M. and A. C. Strauss (2007). *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory* (3rd ed.). Sage Publications, Inc.
- Crowston, K., H. Annabi, and J. Howison (2003). Defining open source software project success. In *Proceedings of International Conference on Information Systems*.
- Crowston, K., K. Wei, J. Howison, and A. Wiggins (2010). Free/libre open source software: What we know and what we do not know. *ACM Computing Surveys* 44(2), 2012. Working Paper.
- de Tocqueville, A. (2004). *Democracy in America*. New York, NY: Library of America.
- Dellarocas, C. (2006). Reputation mechanisms. In *Handbook on Economics and Information Systems*, pp. 2006. Elsevier Publishing.
- Dellarocas, C. (2010, June). Designing reputation systems for the social web. *SSRN eLibrary*.

- DiPrete, T. A. and G. M. Eirich (2006). Cumulative advantage as a mechanism for inequality: A review of theoretical and empirical developments. *Annual Review of Sociology* 32(1), 271–297.
- Earl, J. and K. Kimport (2011). *Digitally Enabled Social Change: Activism in the Internet Age*. MIT Press.
- Eisenhardt, K. M. (1989, October). Building theories from case study research. *The Academy of Management Review* 14(4), 532–550.
- English, R. and C. M. Schweik (2007). Identifying success and tragedy of FLOSS commons: A preliminary classification of sourceforge.net projects. In *Proceedings of the First International Workshop on Emerging Trends in FLOSS Research and Development*, pp. 11. IEEE Computer Society.
- Forte, A., V. Larco, and A. Bruckman (2009). Decentralization in wikipedia governance. *Journal of Management of Information Systems* 26(1), 49–72.
- Freeman, J. (1973). The tyranny of structurelessness. *Berkeley Journal of Sociology* 17, 151–165.
- Fuster Morell, M. (2010). *Governance of Online Creation Communities: Provision of infrastructure for the building of digital commons* (Ph.D. Thesis ed.). European University Institute. 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.
- Gini, C. (1997). Concentration and dependency ratios. *Rivista di Politica Economica* 87, 769–792.
- Goffman, E. (1974). *Frame Analysis: An Essay on the Organization of Experience* (First Edition ed.). Harvard University Press.
- Halfaker, A., A. Kittur, and J. Riedl (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*, New York, NY, USA, pp. 163–172. ACM.
- Hargadon, A. B. and Y. Douglas (2001, September). When innovations meet institutions: Edison and the design of the electric light. *Administrative Science Quarterly* 46(3), 476–501.
- Healy, K. and A. Schussman (2003). The ecology of open-source software development.
- Hess, C. and E. Ostrom (Eds.) (2011). *Understanding Knowledge as a Commons: From Theory to Practice*. The MIT Press.

- Hill, B. M. (2005, October). Reflections on free software past and future. *First Monday* 10(10). Special issue.
- Hill, B. M., C. Burger, J. Jesse, and J. Bacon (2008). *Official Ubuntu Book, The* (3 ed.). Prentice Hall PTR.
- Hill, B. M. and A. Monroy-Hernández (2013, May). The remixing dilemma the trade-off between generativity and originality. *American Behavioral Scientist* 57(5), 643–663.
- Hill, B. M., A. Monroy-Hernández, and K. Olson (2010). Responses to remixing on a social media sharing website. In *Proceedings of the 4th AAAI Conference on Webslogs and Social Media*, Washington, D.C., pp. 74–81. AAAI.
- Howe, J. (2008, August). *Crowdsourcing: Why the power of the crowd Is driving the future of business* (1 ed.). Crown Business.
- Huang, R. (2012). RQDA: r-based qualitative data analysis. R package version 0.2-3.
- Ingo, H. (2006, August). *Open Life: The Philosophy of Open Source*. Lulu.com.
- Ito, M., D. Okabe, and I. Tsuji (Eds.) (2012). *Fandom Unbound: Otaku Culture in a Connected World*. Yale University Press.
- Jaszi, P. (1991, April). Toward a theory of copyright: The metamorphoses of "Authorship". *Duke Law Journal* 1991(2), 455.
- Jenkins, H. (2008). *Convergence Culture: Where Old and New Media Collide* (Revised ed.). New York: NYU Press.
- Jenkins, H., K. Clinton, R. Purushotma, A. Robinson, and M. Weigel (2006, October). Confronting the challenges of participatory culture: Media education for the 21st century. Technical report, MacArthur Foundation, Chicago, Illinois, USA.
- Jenkins, J. C. (1977). Radical transformation of organizational goals. *Administrative Science Quarterly* 22(4), 568–586.
- Karpf, D. (2012). *The MoveOn Effect: The Unexpected Transformation of American Political Advocacy* (First Edition ed.). Oxford University Press, USA.
- Keen, A. (2007). *The Cult of the Amateur: How Today's Internet is Killing Our Culture* (3rd Printing ed.). Crown Business.
- Kittur, A., E. Chi, B. Suh, and T. Mytkowicz (2007). Power of the few vs. wisdom of the crowd: Wikipedia and the rise of the bourgeoisie. San Jose, California.
- Kittur, A. and R. E. Kraut (2010). Beyond wikipedia: coordination and conflict in online production groups. In *Proceedings of the 2010 ACM conference on Computer supported cooperative work*, Savannah, Georgia, USA, pp. 215–224. ACM.

- Kollock, P. (1999). The economies of online cooperation: Gifts and public goods in cyberspace. In M. Smith and P. Kollock (Eds.), *Communities in Cyberspace*, pp. 220–239. London: Routledge.
- Konieczny, P. (2009). Governance, organization, and democracy on the internet: The iron law and the evolution of wikipedia. *Sociological Forum* 24(1), 162–192.
- Kraut, R. E. and P. Resnick (2012, March). *Building Successful Online Communities: Evidence-Based Social Design*. The MIT Press.
- Krishnamurthy, S. (2002, June). Cave or community? an empirical examination of 100 mature open source projects. *First Monday* 7(10).
- Krishnamurthy, S. (2005, October). The elephant and the blind men - deciphering the Free/Libre/Open source puzzle. *First Monday* 10(10). Special issue.
- Lakhani, K. R. and E. von Hippel (2003, June). How open source software works: "Free" user-to-user assistance. *Research Policy* 32(6), 923–943.
- Lave, J. and E. Wenger (1991). *Situated learning: Legitimate peripheral participation*. Cambridge University Press.
- Leach, D. K. (2005). The iron law of what again? conceptualizing oligarchy across organizational forms. *Sociological Theory* 23(3), 312–337.
- Lerner, J. and J. Tirole (2002). Some simple economics of open source. *Journal of Industrial Economics* 50(2), 197–234.
- Lessig, L. (2008). *Remix: Making Art and Commerce Thrive in the Hybrid Economy*. Penguin Press HC.
- Leuf, B. and W. Cunningham (2001). *The Wiki way: quick collaboration on the Web*. Addison-Wesley Longman Publishing Co., Inc. Boston, MA, USA.
- Levenshtein, V. I. (1966). Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics Doklady* 10(8), 707–710.
- Lewis, L. A. (Ed.) (1992). *The Adoring Audience: Fan Culture and Popular Media*. Routledge.
- Lih, A. (2009). *Wikipedia Revolution, The: How a Bunch of Nobodies Created the World's Greatest Encyclopedia* (First Edition. First Printing. ed.). Hyperion.
- Lin, Y. (2005, October). The future of sociology of FLOSS. *First Monday* 10(10). Special issue.
- Lipset, S. M., M. A. Trow, and J. S. Coleman (1956). *Union democracy: The Internal Politics of the International Typographical Union*. Glencoe, Ill: Free Press.

- Loubser, M. (2010, July). *Organisational Mechanisms in Peer Production: the Case of Wikipedia*. DPhil, Oxford Internet Institute, Oxford University, Oxford, UK.
- Loveland, J. and J. Reagle (2013). Wikipedia and encyclopedic production. *New Media & Society*.
- Luther, K., K. Caine, K. Zeigler, and A. Bruckman (2010). Why it works (when it works): Success factors in online creative collaboration. In *Proceedings of the ACM Conference on Supporting Group Work.*, Sanibel Island, Florida, USA. ACM.
- MacCormack, A., J. Rusnak, and C. Y. Baldwin (2006). Exploring the structure of complex software designs: An empirical study of open source and proprietary code. *Management Science* 52(7), 1015–1030.
- MacCormack, A. D., J. Rusnak, and C. Y. Baldwin (2012, October). Exploring the duality between product and organizational architectures: A test of the “mirroring” hypothesis. *Research Policy* 41(8), 1309–1324.
- Manovich, L. (2005). Remixing and remixability.
- McCarthy, J. D. and M. N. Zald (1977, May). Resource mobilization and social movements: A partial theory. *The American Journal of Sociology* 82(6), 1212–1241.
- Michels, R. (1915). *Political Parties: A Sociological Study of the Oligarchical Tendencies of Modern Democracy*. New York, NY: Hearst International Library Co.
- Minkoff, D. C. (1999, May). Bending with the wind: Strategic change and adaptation by women’s and racial minority organizations. *The American Journal of Sociology* 104(6), 1666–1703.
- Mockus, A., R. T. Fielding, and J. D. Herbsleb (2002). Two case studies of open source software development: Apache and mozilla. *ACM Transactions on Software Engineering Methodology* 11(3), 309–346.
- Monroy-Hernández, A., B. M. Hill, J. Gonzalez-Rivero, and d. boyd (2011). Computers can’t give credit: how automatic attribution falls short in an online remixing community. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, Vancouver, BC, Canada, pp. 3421–3430. ACM.
- Murray, F. and S. O’Mahony (2007, November). Exploring the foundations of cumulative innovation: Implications for organization science. *Organization Science* 18(6), 1006–1021.
- Nelson, T. H. (1981). *Literary Machines* (3rd ed.). Swarthmore, PA: Self Published.
- Nielsen, J. (1993). *Usability Engineering*. Academic Press Inc.
- O’Connor, S. D. (1991). Feist publications, inc. v. rural telephone service co.

- O'Mahony, S. and F. Ferraro (2007). The emergence of governance in an open source community. *Academy of Management Journal* 50(5), 1079–1106.
- Oostendorp, N. (1999, November). Everything.
- Orlikowski, W. J. (1992). Learning from notes: organizational issues in groupware implementation. In *Proceedings of the 1992 ACM conference on Computer-supported cooperative work*, Toronto, Ontario, Canada, pp. 362–369. ACM.
- Orlikowski, W. J. and D. C. Gash (1994). Technological frames: making sense of information technology in organizations. *ACM Transactions of Information Systems* 12(2), 174–207.
- Ortega, F. (2009, April). *Wikipedia: A Quantitative Analysis*. PhD dissertation, Universidad Rey Juan Carlos.
- Palfrey, J. and U. Gasser (2008). *Born Digital: Understanding the First Generation of Digital Natives*. Basic Books.
- Papert, S. (1980). *Mindstorms: Children, computers, and powerful ideas*. New York: Basic Books.
- Piven, F. F. and R. A. Cloward (1977). *Poor people's movements: why they succeed, how they fail* (1st ed ed.). New York: Pantheon Books.
- Polletta, F. (2013). Participatory democracy in the new millennium. *Contemporary Sociology: A Journal of Reviews* 42(1), 40–50.
- Postma, R. (2000). Everything FAQ.
- Preece, J. (2000). *Online Communities: Designing Usability and Supporting Sociability* (1 ed.). Wiley.
- Raymond, E. S. (1999). *The cathedral and the bazaar: Musings on Linux and open source by an accidental revolutionary*. Sebastopol, CA: O'Reilly and Associates.
- Reagle, J. (2009). Wikipedia: the happy accident. *interactions* 16(3), 42–45.
- Reagle, J. M. (2008). *In Good Faith: Wikipedia Collaboration and the Pursuit of the Universal Encyclopedia*. PhD, New York University.
- Reich, J., R. Murnane, and J. Willett (2012, January). The state of wiki usage in U.S. k–12 schools leveraging web 2.0 data warehouses to assess quality and equity in online learning environments. *Educational Researcher* 41(1), 7–15.
- Resnick, M., J. Maloney, A. Monroy-Hernández, N. Rusk, E. Eastmond, K. Brennan, A. Millner, E. Rosenbaum, J. Silver, B. Silverman, and Y. Kafai (2009). Scratch: programming for all. *Commun. ACM* 52(11), 60–67.

- Rucht, D. (1999). Linking organization and mobilization: Michels's iron law of oligarchy reconsidered. *Mobilization: An International Quarterly* 4(2), 151–169.
- Salganik, M. J., P. S. Dodds, and D. J. Watts (2006, February). Experimental study of inequality and unpredictability in an artificial cultural market. *Science* 311(5762), 854–856.
- Schweik, C. and R. English (2012). *Internet Success: A Study of Open Source Software Commons*. Cambridge, Massachusetts: MIT Press.
- 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.
- Shirky, C. (2010). *Cognitive Surplus: Creativity and Generosity in a Connected Age*. Penguin Press HC, The.
- Singer, J. D. and J. B. Willett (2003). *Applied Longitudinal Data Analysis: Modeling Change and Event Occurrence* (1 ed.). Oxford University Press, USA.
- Sinnreich, A. (2010). *Mashed Up: Music, Technology, and the Rise of Configurable Culture*. University of Massachusetts Press.
- Snow, D. A., E. B. Rochford, S. K. Worden, and R. D. Benford (1986, August). Frame alignment processes, micromobilization, and movement participation. *American Sociological Review* 51(4), 464–481.
- Soule, S. A. and B. G. King (2008, May). Competition and resource partitioning in three social movement industries. *The American Journal of Sociology* 113(6), 1568–1610.
- Stallman, R. M. (1999). Free universal encyclopedia and learning resource.
- Stallman, R. M. (2002, October). The GNU manifesto. In *Free Software, Free Society: Selected Essays of Richard M. Stallman*, pp. 224. Free Software Foundation.
- Tapscott, D. and A. D. Williams (2008). *Wikinomics: How Mass Collaboration Changes Everything* (Expanded ed.). Portfolio Hardcover.
- Thom-Santelli, J., D. R. Cosley, and G. Gay (2009). What's mine is mine: Territoriality in collaborative authoring. In *Proceedings of the 27th international conference on Human factors in computing systems*, Boston, MA, USA, pp. 1481–1484. ACM.
- Tufekci, Z. and C. Wilson (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., M. Wattenberg, M. M. McKeon, and D. Schuler (2007). The hidden order of wikipedia. In *Online Communities and Social Computing*, HCII, pp. 445–454. Berlin, Heidelberg: Springer-Verlag.
- von Hippel, E. (2005). *Democratizing innovation*. Cambridge, Massachusetts: The MIT Press.
- von Hippel, E. and G. von Krogh (2003, March). Open source software and the 'private-collective' innovation model: Issues for organization science. *Organization Science* 14(2), 209–223.
- von Krogh, G. and E. von Hippel (2006, July). The promise of research on open source software. *Management Science* 52(7), 975–983.
- Voss, K. and R. Sherman (2000). Breaking the iron law of oligarchy: Union revitalization in the american labor movement. *American Journal of Sociology* 106(2), 303–349.
- Walston, C. E. and C. P. Felix (1977). A method of programming measurement and estimation. *IBM Systems Journal* 16(1), 54–73.
- Watts, D. J. (2011). *Everything Is Obvious: How Common Sense Fails Us*. Crown Business.
- Weber, S. (2004). *The Success of Open Source*. Harvard University Press.
- Wesler, H. T., M. Smith, D. Fisher, and E. Gleave (2008). Distilling digital traces: Computational social science approaches to studying the internet. In *The SAGE Handbook of Online Research Methods*, pp. 116–140.
- West, J. and S. O'Mahony (2005). Contrasting community building in sponsored and community founded open source projects. In *Proceedings of the 38th Annual Hawaii International Conference on System Sciences (HICSS'05) - Track 7 - Volume 07*, pp. 196.3. IEEE Computer Society.
- Willer, R. (2009a, February). Groups reward individual sacrifice: The status solution to the collective action problem. *American Sociological Review* 74(1), 23 –43.
- Willer, R. (2009b). A status theory of collective action. In *Advances in Group Processes*, Volume 26, pp. 133–163. Bingley: Emerald Group Publishing.
- Williamson, O. E. (1981, November). The economics of organization: The transaction cost approach. *The American Journal of Sociology* 87(3), 548–577.
- Wu, F., D. M. Wilkinson, and B. A. Huberman (2009). Feedback loops of attention in peer production. In *Proceedings of the 2009 International Conference on Computational Science and Engineering - Volume 04, CSE '09*, Washington, DC, USA, pp. 409–415. IEEE Computer Society.

Zhu, H., R. Kraut, and A. Kittur (2012). Organizing without formal organization: group identification, goal setting and social modeling in directing online production. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work, CSCW '12*, New York, NY, USA, pp. 935–944. ACM.

Zhu, H., R. E. Kraut, Y.-C. Wang, and A. Kittur (2011). Identifying shared leadership in wikipedia. In *Proceedings of the 2011 annual conference on Human factors in computing systems, CHI '11*, New York, NY, USA, pp. 3431–3434. ACM.

Zittrain, J. (2008). *The Future of the Internet—And How to Stop It*. Yale University Press.

Zuckerman, E. W. (1999, March). The categorical imperative: Securities analysts and the illegitimacy discount. *The American Journal of Sociology* 104(5), 1398–1438.

Acknowledgments

This dissertation is about cooperation but it is also the product of it. I have deep gratitude for the many people who have helped me build, and build toward, this dissertation over the last five years.

In the most direct sense, this includes my co-authors on this work: Andrés Monroy Hernández contributed to *The Remixing Dilemma* and Aaron Shaw contributed to *Laboratories of Oligarchy* and gave me detailed feedback on my broader research agenda. Both Aaron and Andrés are collaborators on other papers and, through working together over the last several years, each has become one of my very dearest friends. Both deserve enormous credit for putting up with me and all my quirks. I'm looking forward to each of them putting up with lots more of me in the future. If all goes well, I will be collaborating with each for many more years.

I also need to thank my thesis committee and program committee (the latter is a subset of the former). I was reminded frequently on the job market that my committee is an impressive bunch. Indeed, every one of them has impressed me since I was in college. If you had told me ten years ago that I would graduate from an MIT with this group as my committee, I would have refused to believe it as simply too good to be true. Most impressive to me, however, is that in addition to their deep intellectual contributions and their role shaping – and even defining – my field, they are all friendly, generous, constructive, and good people. It's not always, maybe even often, like that for other students. I could not have wished for a better group.

I would not have done a PhD at all if Eric von Hippel had not found me at a lecture in Jonathan Zittrain's class, invited me to dinner, and convinced to come back to MIT.

Eric provided me with the opportunity, the guidance, and the support to succeed at every stage of my PhD. I know of no other faculty members who gave their students more, and went to bat for their students more consistently, than Eric did for me. He challenged me to defend my ideas and taught me the importance of taste and strategy. One long and challenging conversation with him in my second year of my PhD was the most important turning point in my intellectual and academic life. He has provided advice, and a template, for what I view as a successful academic career. I hope I can accomplish a tiny fraction of what he has already achieved in his own. I can't express my gratitude.

Up the river, Yochai Benkler has provided me with a home away from MIT. Yochai invited me in as a collaborator on a paper and he's let me stay as the scope of our work together expands. He has supported me financially, helped arrange for my position as a Berkman Fellow, and provided the impetus for the Berkman Cooperation group. He is too generous with his time and ideas and has given me more than he ever needed to. Without Yochai's idea of "peer production," I wouldn't have the conceptual tools to tie these three essays into a coherent whole. I will continue to look up to Yochai and feel lucky to be able to work with him and learn from him.

After going through two Media Lab advisors in three years, Mitch provided me with the supervisor in the Media Lab I had been looking for. Mitch was supportive, intellectually engaging, and helped me thrive. Mitch has welcomed me as a peripheral member of the Scratch research team and provided me with access to the data I used in a majority of my current publications. He has given me detailed advice on every piece of research I've done with Scratch, at multiple stages of development, and has improved each piece enormously. His students and group are some of my best friends at the lab and I feel extremely lucky to count him as one as well.

I met with Tom Malone before I started graduate school and his work was an important driver of my decision to start this dissertation. His advice and feedback, and support from him and from the wonderful group he has put together at the Center for Collective Intelligence, has been a critical part of my development at MIT. During my first year, he came every week to a small reading group I put together on collective intelligence which

framed the way I address the question in this dissertation. Tom's contributions to my work cannot be overvalued.

Because four is already a large committee, I was not able to ask Susan Silbey, Ezra Zuckerman, Roberto Fernandez, or John Willett to be on my committee. But I *really* wanted to ask each of them. I think of each as mentors and feel that each have driven the way that I think about research. I see their fingerprints all over the better parts of my scholarship and I hope they do too. I cannot thank them enough for all they have given me. Susan, in particular, advised my work on *Almost Wikipedia* directly and it would not have been possible without her training, her supervision, and her generous advice.

My intellectual home over the last two years was really centered in the Berkman Cooperation group. In addition to Aaron and Andrés, it's worth singling out Mayo Fuster Morell, Jérôme Hergueux, Dariusz Jemielniak, Brian Keegan, and Charlie DeTar. I hope to get the band back together for a reunion tour sometime. Thanks to Becca Tabasky for making it all possible and for being such a good friend on top of it.

Of course, there is a huge community of other students who supported me deeply and who helped me flounder my way toward a solid set of papers. My cohort in the Sloan PhD program, and in the rest of the TIES group, helped me enormously along the way. My officemates Ben Rissing and Heekyung Hellen Kim were generous and useful sounding boards as were Abhishek Nagaraj, Ethan Mollick, Michael Bikard, Chuck Eesley, and John-Paul Ferguson. Karim Lakhani graduated before I arrived but helped enormously nonetheless. In the Media Lab, I was helped enormously by the Scratch team and especially by Amos Blanton, Ricarose Roque, and Sayamindu Dasgupta. I was also given an intellectual home and affiliation during the first years of my PhD at the Center for (Future?) Civic Media. Thanks in particular to Chris Csikszentmihályi, Ethan Zuckerman, Lorrie LeJeune, Pablo Rey, Charlie DeTar, and J. Nathan Matias. I've still got my books in your offices. Thanks for that and for everything else.

Also thanks to advisors of my previous work that set the stage and provided the tools for this dissertation. This includes Walter Bender, Biella Coleman, Chris Csikszentmihályi, Jim Miller, and Herb Bernstein. Leah Buechley was also a wonderful mentor and

collaborator. Thanks to Paul Sawaya for volunteering to write the software that collected wiki data. If he had not stepped up, the third essay in this dissertation could not have happened.

I have presented work in all three essays in a variety of settings and received valuable feedback which has shaped and improved this work. For *Almost Wikipedia*, I am grateful to the Berkman Center Luncheon Series, the Digital Communities conference, the Center for Collective Intelligence, the MIT Economic Sociology Working Group (ESWG), the CyberScholars MetaForum, the Open and User Innovation Workshop, the TIES seminar at Sloan, the Social Media Collective at Microsoft Research in Cambridge, Wikimania, and the Wikimedia Foundation.

For *The Remixing Dilemma*, I must thank Andrea Forte and Cliff Lampe and several anonymous reviewers for their valuable feedback on the piece while it was under review and revision at the *American Behavioral Scientist*. I must also acknowledge the Berkman Cooperation Group, Wikipedia Academy, Michigan MUG, and the University of Washington Department of Communication for their feedback on earlier drafts of this work.

Finally, for *Laboratories of Oligarchy*, I must thank the Berkman Cooperation Group, Lance Bennett and Helen Margetts, and the whole “online collective action” ECPR Joint Sessions working group, the Princeton Center for Information Technology Policy, and the ESWG for their reactions to presentations of earlier version of the work.

Thanks also to my housemates at the Acetarium and at Mirthaterra including Erik Garrison, Adelaida McIntire, Dafydd Harries and, of course, all the wonderful Acetarium Residents over the years. In terms of this document, Erik Garrison deserves particular recognition for the work he did to write the software I used to parse the wikis in the final essay. Bernie Innocenti and SJ Klein each argued with me over all of this work and I thanks them for their contributions as well. Thanks as well go to the Boston free software community including #acetarium, #yukidoke, the staff and management at the Free Software Foundation, and Iron Bloggers.

Of course, there would be no research, and no point, if not for all the people who

contribute to peer production. That includes all the mostly anonymous contributors to Scratch, to Wikia wikis, and all the people who let me interview them for *Almost Wikipedia*. Beyond their contributions to my research, their time and effort means the world is a slightly better place than it would be otherwise. We all owe them.

And *thank you, thank you, thank you* to all the wonderful people who have helped me enormously and who I did not mention here. Hopefully, you all know me well enough to know I'm super bad at remembering this kind of thing. Sorry! Of course, thanks are due to my parents, Carter Hill and Winnie Mann, who gave me the tools, opportunity, and encouragement. Most of all, thanks to my partner Mika Matsuzaki. She's made this possible. It will be her turn soon.

Appendix A

Almost Wikipedia Interview Protocol

Thank you so much for taking the time to talk to me!

I'm a PhD student at MIT doing research on online encyclopedia projects and I am interested in [PROJECT NAME] which I know you were involved in starting. I've gone through some materials online that I dug up about [PROJECT NAME] and I'm interested in learning more and hearing the whole story from your perspective.

I've got a series of open-ended questions written down but you should feel welcome to take the conversation in directions you think are interesting. We can keep this pretty unstructured. Please let me know, at any point, if you have any questions for me and don't feel like you have to answer any questions you don't want to.

My plan is to write this up into a academic paper. If you'd prefer, I'll be happy to disguise your name and project in the published version and to keep my notes and recordings confidential. You can let me know later as well.

Background

First, I'd love for you to tell me a little about yourself today – where you are and what you're up to?

Where do you live? *(Most of these interviews will be phone based.)*

If you don't mind me asking, how old are you now?

Can you tell me a little bit about your passions or hobbies?

Thinking back in time to when you created [PROJECT NAME], can you tell me about what you were up to then?

Where were you living?

How were you supporting yourself? Were you working?

If you don't mind my asking, how do you support yourself now?

How about hobbies or passions at the time?

Project Creation

When did you first decide to create an online encyclopedia? What were the there projects or experiences that led up your decision to start [PROJECT NAME]?

Did you work with others to create [PROJECT NAME]? Who? What was the collaboration like?

Were you influenced by other projects, thinkers, or writers? Who/which ones? How?

Why did you start the project? What were the ideas or goals behind [PROJECT NAME]?

Can you describe the basic “design” of [PROJECT NAME]? How did it work? What did it look like?

Who did you imagine as the audience for the [PROJECT NAME]? What type of people were you trying to serve?

How did you try to market your project?

Do you think you were successful? Why? Or, why not?

Who did you imagine as the contributors to [PROJECT NAME]? What type of people did you imagine contributing?

How did you try to recruit contributors?

Do you think you were successful? Why? Or, why not?

What sort of resources were necessary (e.g., money, technical skills, labor, etc.)? What did you do to get these? Did you have trouble getting necessary resources?

At the time of your projects, there were other online collaborative encyclopedia projects. Did you know about these? What did/do you think of them?

[*Going through my list.*] Did you know about [OTHER PROJECT NAME]? What did you think of it? How was your project different? What did you think of them?

Are there other similar projects you think I’ve missed?

About Your Project

Can you walk me through the history of the [PROJECT NAME]?

What were some of the greatest challenges you faced with [PROJECT NAME]?

What are some of your proudest achievements with [PROJECT NAME]? What did you think you did right, and why?

What are some of your greatest regrets? What do you think you did wrong, and why?

Your Project Today

Where, if anywhere, is [PROJECT NAME] today?

If you could do your project over again, what would you do differently?

More generally, what has your experience with [PROJECT NAME] taught you? What have others learned?

Are you still in touch with any collaborators? What are they up to?

Wikipedia

Superficially at least, Wikipedia seems to share a lot in common with [PROJECT NAME]. What do you think about Wikipedia?

How is Wikipedia similar, and different, to [PROJECT NAME]?

What has it been like watching Wikipedia become so enormously successful?

Have you contributed to Wikipedia? How do think the experience compared to work on [PROJECT NAME]? What was different?

Why do you think Wikipedia has worked out so much better than [PROJECT NAME]?

What do you think Wikipedia could learn from [PROJECT NAME]?

Do you think you've learned anything about [PROJECT NAME] from Wikipedia?

Concluding

Thank you *so much* for taking the time to talk to me! I appreciate all the information you've been able to give me.

If you have any questions for me, I'm happy to answer them. If you're curious about my research project or my hypotheses, I'm happy to give you details.

If you know of other people I might want to talk to or other ideas of projects I might have missed, I'd love to hear that.