Journal of Information Technology & Politics
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/witp20

Open Source Political Community Development: A Five-Stage Adoption Process

David Karpf [a b]
[a] Rutgers University
[b] Eagleton Institute of Politics

Available online: 01 Jun 2011

To cite this article: David Karpf (2011): Open Source Political Community Development: A Five-Stage Adoption Process, Journal of Information Technology & Politics, 8:3, 323-345

To link to this article: http://dx.doi.org/10.1080/19331681.2011.575020

Journal of Information Technology & Politics, 8:323–345, 2011

ISSN: 1933-1681 print/1933-169X online
DOI: 10.1080/19331681.2011.575020

Routledge
Taylor & Francis Group

Open Source Political Community Development: A Five-Stage Adoption Process

David Karpf

ABSTRACT. This article considers the emergence of large-scale "commons-based peer production" projects such as Wikipedia.org from an institutional development perspective. The argument it makes is threefold. First, that that the lowered transaction costs and information abundance found online transform a subset of public goods problems, essentially replacing free-ridership with mass coordination as the central challenge. Second, that the boundaries of this subset are defined by a "power-law topology" that leads to the emergence of online hub spaces and serves to resolve search problems endemic to the anti-geographic online landscape. These boundary conditions limit the overall impact of commons-based peer production for the political space. Third, that all such hubs move through a common five-stage institutional development process, directly related to standard models of the diffusion of innovation. Identification of the institutional development process behind Wikipedia leads in turn to the stipulation of seven hypotheses: the "Field of Dreams" fallacy, the "Interest Horizons" thesis, "Political Strategy Is Not Like Computer Code," the "Location-based Wave" thesis, "Power Law Fragility Under Moore's Law," the "Punctuated Equilibrium" thesis, and "Code-Forking the Public Sphere." Each thesis holds direct implications for the potential and limitations of "open source" applications in the political arena.

KEYWORDS. online communities, open source, Wikipedia

> Open source is not a piece of software, and it is not unique to a group of hackers. Open source is a way of organizing production, of making things jointly.
> —Steven Weber, *The Success of Open Source*

> Unlike previous reference works that stand on library shelves distanced from the institutions, people, and discussions from which they arose, Wikipedia is both a community and an encyclopedia.
> —Joseph Reagle, *Good Faith Collaboration*

"Open Source" refers to a type software, to the process through which that software is produced, *and* to the online community that is

Dave Karpf (PhD, Political Science, University of Pennsylvania) is an assistant professor of journalism and media studies at Rutgers University, as well as a faculty associate at the Eagleton Institute of Politics. His research concerns the impact of the Internet on American political associations, including community blogs and Internet-mediated advocacy groups.

The research presented here was conducted while Dr. Karpf was a postdoctoral research associate at Brown University's Taubman Center for Public Policy. He would like to thank the participants at the *JITP* 2010 Politics of Open Source conference for helpful feedback, as well as Rogers Smith, Jack Nagel, John Lapinski, Lokman Tsui, Daniel Kreiss, and Rasmus Kleis Nielsen for substantial feedback on the piece.

Address correspondence to: David Karpf, Rutgers University, 4 Huntington St., New Brunswick, NJ 08901 (E-mail: dkarpf@rutgers.edu).

responsible for it. The software has attracted much attention thanks in particular to the successes of Linux and Apache—sophisticated, broadly adopted software programs with large, complex open source communities behind them. For 15 years, since the demonstrated success of Linux, many have asked "What other elements of society could benefit from open source as a process or be transformed by open source–type communities? Can we have open source journalism, open source culture, open source politics?" Benkler (2006) has termed the process as "commons-based peer production" (p. 59) and contrasts it to markets and firms. Shirky (2008) focuses more on the community, discussing how the lowered transaction costs of the Internet enable "organizing without organizations." Bruns (2008) suggests "produsage" as a term-of-art for the new, Internet-mediated production processes. Kelty (2008) offers an ethnographic account of the open source software community, describing it as a "recursive public" that deliberates both through words and through the production of software code. All of these authors highlight the surprising success of open source as a product, a process, and a community, and help illuminate how and why it functions the way it does.

It is abundantly clear from these authors that, when a massive community devoted to commons-based peer production emerges, it is capable of producing sophisticated goods that manage to overcome the traditional collective action problem (Olson, 1965). But we continue to know relatively little about the conditions under which such communities manage to form. Open source solutions do not work for every problem, nor do robust open source communities form out of thin air. This article treats the development of online "collaborative communities" (Reagle, 2010, p. ix), through a diffusion and institutional development lens. It asks how online communities develop, and argues that a five-stage development process, tied to the standard diffusion curve, is likely a central feature of their development. Focusing on the sequence of institutional development challenges that face a developing collaborative community is particularly helpful in clarifying the boundaries of open source–style activity in the political arena.

Particularly after the Dean campaign of 2003–2004, a set of scholars, public intellectuals, and practitioners has shown strong interest in the potential emergence of "open source politics" (Bruns, 2008; Fine, Sifry, Raseij, & Levy, 2008; Hara, 2008; Jenkins, 2006; Lessig, 2003; Ratcliffe & Lebkowsky, 2005; Rushkoff, 2003; Sifry, 2004; Trippi, 2004). Steven B. Johnson, for instance, suggests that, "Using open-source coding as a model, it's not a stretch to believe the same process could make politics more representative and fair. Imagine, for example, how a grassroots network could take over some of the duties normally performed by high-priced consultants who try to shape a campaign message that's appealing" (Sifry, 2004, p. 6). The use of the Internet to spur citizen participation has become a growth industry, fueling dozens of political consultancies and popular books. Lost underneath is any clear understanding of what, specifically, open source politics is likely to entail. Hindman (2007), for instance, cautions that we are indeed entering an era of open source politics, but reminds us that open source production features the replication of the same elite systems that have long dominated politics. We can define "open source politics" as the application of Internet-mediated collaborative tools by a motivated community-of-interest to collective action problems that previously were solved exclusively by markets or firms. The question that then follows is, under what conditions will these communities flourish, and under what conditions will they fail?

This article offers an analysis of the online community-formation process, treating Wikipedia as a guiding case example and drawing upon the robust literature that has emerged in recent years around Wikipedia as an "impossible good" (Ciffolilli, 2003; Kollock, 1999). By understanding the technological affordances that make Wikipedia possible, and by understanding the community formation pattern that Wikipedia and other online communities experience as they endure a *diffusion process*, it allows us to postulate a five-stage institutional development process and leads to seven hypotheses that clarify the potential and limitations of open source-type community participation in the political arena.

WHY STUDY WIKIPEDIA?

Wikipedia, "the encyclopedia that anyone can edit," is a public good, both in definition and in spirit. In the success of Wikipedia, there lies a significant puzzle for social scientists. Much of the foundational literature on collective action and public goods problems suggests that Wikipedia *should not* exist. Free ridership should prevent it from forming. As a public good—being both nonrival (my use of Wikipedia does not reduce your ability to use it) and nonexclusive (all people can access Wikipedia, regardless of whether they contributed to it)—Wikipedia ought to suffer from underprovision. Yet, as the seventh-most-visited site on the entire World Wide Web,[1] and with over 10 million articles across 260 languages (Reagle, 2010), the one critique that the site has never been subjected to is "there just isn't enough of it."

Though not a *political* endeavor itself, Wikipedia provides a mature example of the type of "open source–like" online communities-of-interest enabled by the structure and novel attributes of Web-based communication protocols. Early critics of Wikipeida noted the crucial difference between an online collaborative encyclopedia and online collaborative software: "Bad, incorrect code doesn't compile. Bad, incorrect information on the 'net lives on and non-experts hardly ever notice the mistake" (Reagle, 2010, p. 78). Stalder has described this as a key difference between "open source" and "open culture" (Stalder, 2006). Given that political ideas, strategies, and decisions also "don't compile," Wikipedia is a more apt comparison to the political arena than Linux, Apache, or other major software projects.

Careful study of such a case can help us to identify the parameters within which such communities function, in turn leading to clearer thinking about the potential of "open source politics." Most importantly, Wikipedia demonstrates that the dramatic reduction in the costs of online communication produces a condition of *information abundance* in which the challenge to mass collaboration approximates a coordination game, rather than a free rider problem. The problem of mass coordination is solved through the development of a *power-law topology*, in which large hub spaces let communities-of-interest engage in collaborative efforts that would have been impossible under previous information regimes. Novel solutions to the collective action problem, and novel structures for collective action, become possible online because of these changes to the costs, abundance, and availability of information. The emerging landscape of collective action fundamentally departs from previous eras in direct relation to the salience of these attributes, and Wikipedia serves as a guiding example for understanding them.

Much of what makes Wikipedia worthy of analysis is the sheer scale of its success. The only Web sites more frequently visited than Wikipedia are the search engines/e-mail providers Google, Yahoo, Windows Live/MSN, video sharing site YouTube.com, and social networking site Facebook.com.[2] Unlike these for-profit Internet giants (YouTube being a subsidiary of Google), Wikipedia operates with a budget of less than $500,000 and a staff of fewer than a dozen employees (Lih, 2009, p. 4). A core of 75,000 active volunteer "Wikipedians," along with 13.3 million registered users and an untold number of anonymous users, contribute the bulk of content, citations, and edits.[3] Despite providing little opportunity for fame or recognition to the volunteers, this expanding volunteer corps has remained resilient against attempts to subvert the site through vandalism or thinly veiled advertising. A 2007 study by the Pew Internet and American Life Project found that 36 percent of American adult Internet users consult Wikipedia. It is especially popular among the well-educated, with 50 percent of all online Americans who hold a college degree using it as a reference. On a typical day, 8 percent of online Americans consult Wikipedia, making it more popular than online purchasing, dating Web sites, setting travel reservations, using chat rooms, and participating in online auctions (Rainie & Tancer, 2007). This raises the practical question of whether Wikipedia is now so distinct as to be something different, something greater, than an encyclopedia. Consider the

following: In the pre-Internet era, what percent of Americans would we suspect consulted an encyclopedia on an average day? Is it likely that looking things up in the *Encyclopedia Brittanica* has ever been more popular in daily life than purchasing goods or trying to find a date?

Political scientists have largely overlooked Wikipedia, though a small group of scholars has begun to look into the more generic impacts of the Internet's lowered communication costs on online collective action. Lupia and Sin (2003), for instance, argue that Mancur Olson's (1965) work is "built from historically uncontroversial assumptions about interpersonal communication. Today, evolving technologies are changing communication dynamics in ways that invalidate some of these once uncontroversial assumptions" (p. 315). They go on to present a formal model that suggests the organizational advantage held by small groups in Olson's day is muted by online communication, while the selective benefits that many groups were once able to offer as an incentive for participation are occasionally undermined by the open access of the Web (disciplinary journals, as one example, now face open-access, free, Web-based competition). Bimber, Flanagin, and Stohl (2005) likewise attempt to reconceptualize collective action as a phenomenon of "boundary crossing" between public and private domains—an indication of how near-costless participation in online petitions and other Web-based pressure tactics has become. Lev-On and Hardin (2008) deal directly with the cases raised by Benkler (Wikipedia included) and offer a theoretical framework for analyzing the "logic of Internet-based collective action" (p. 6), arguing that the lowered transaction costs and formation of online hubs allows success in the face of widespread free ridership. They argue that phenomena such as Wikipedia can be conceptualized within the existing framework for collective action studies. My analysis is similar to theirs, but focuses more centrally on community formation than on production process.

The following section will use the Wikipedia example to synthesize several core concepts regarding Internet-mediated communication. Centrally, it will demonstrate that the structure of the Web supports the development of large-scale communities that, benefiting from strong "network effects," can produce tremendous public goods on the basis of surplus labor contributions from hobbyists and partisans. When the costs of participation approach zero, a more complete demand curve for political engagement is revealed. In so doing, the section clarifies *how* online communication differs from previous communication regimes, and also places focus on *when* and *where* such differences are likely to be present. It is only when online *power-law hubs* successfully develop that these non-market, non-firm-based solutions become viable. The article then turns to the diffusion-of-innovations literature (Rogers, 2003; Von Hippel, 2005), arguing that online community formation follows a similar path, and that each stage of diffusion presents a distinct institutional development challenge that can prevent hub formation. A third section then derives a set of hypotheses regarding the stability of *power-law hubs* over time, and the conclusion then makes several points about the political implications of these hypotheses.

THE SUCCESS OF WIKIPEDIA: EASY, FUN, AND FULL OF NETWORK EFFECTS

Wikipedia was founded in 2001 after Nupedia, an attempt at developing an online encyclopedia based on traditional expert-produced and -reviewed contributions, failed to gather momentum (Reagle, 2010). Jimmy Wales had launched Nupedia as an open-access competitor to pricey encyclopedias like *Britannica*. His expectation was that the speed and ease of e-mail communication could lower the costs of producing a high-quality encyclopedia, making the information free for all visitors.[4] Nupedia was to be expert-led, with a traditional (and daunting) seven-stage peer-review and editing process. What Wales and his collaborator Larry Sanger learned was that the increased speed of e-mail alone does little to transform production processes. The hefty editing process resulted in numerous bottlenecks, leading to an estimated 25 articles in its first three years. As academic journal editors have likewise learned, moving

from the fax, phone, and mail systems to digital communication alleviates *some* elements of peer review and content production, but the overall savings prove marginal.

In attempting to radically simplify the production process, Wales and Sanger turned to the "wiki" (from the Hawaiian word "wikiwiki," translating directly to "fast" or "speedy") software platform. Wiki software code enables open content creation and peer editing. Any user with access (and on Wikipedia, most articles are accessible by all) can click an "edit this" button, make changes to the document, and have those changes instantly available to other users. Past versions are automatically archived and viewable, making such experimentation a low-risk affair.

Developer Larry Sanger wrote a memo to the 2,000-member Nupedia mailing list at the launch of the wiki-based site, saying, "Humor me. Go there and add a little article. It will take all of five or ten minutes" (Shirky, 2008, p. 113). With the bottlenecks eliminated, the community responded, producing over 1,000 articles within the first month, and 10,000 within nine months. Clay Shirky describes this as a general shift enabled by the Internet-based information regime: from "filter, then publish," to "publish, then filter." Print-based publication is costly and (thus) scarce. Firms are necessary to provide editorial and quality-control decisions at the front end, ensuring that the final product is well-written and attractive to a paying audience. Comparatively, Shirky notes that Web-based publication is "ridiculously easy." Indeed, ever-expanding transistor capacity and server space render the Web an abundant information environment where point-to-point communication (e-mail) can happen near instantaneously and self-publication is free. Wikipedia could not exist without Internet-mediated communication, and moreover, it could only exist through the embrace of novel alternatives to traditional production practices. Faster and cheaper communications media alone produce little change, but they create the possibility for novel structures for mass collaboration and collective action.

The *ease* of publishing online may be self-evident, but that is a far cry from assuring high-quality encyclopedia entries. Indeed, Wikipedia's quick rise in popularity was accompanied by an avalanche of skepticism regarding the quality of the new-entrant encyclopedia ("free, and worth every penny," so to speak). Jim Giles published a 2005 study in *Nature* magazine challenging this claim through a comparison of Wikipedia and the *Enclopedia Britannica.* Peer reviewers recruited by *Nature* found an average of four inaccuracies per Wikipedia article, and three per equivalent *Britannica* article (Giles, 2005). A back-and-forth firestorm ensued, with *Britannica* staff criticizing the study and demanding a retraction. *Nature* offered a clarification of its methodology, but stood by the study and refused to retract it. Physicists Huberman and Wilkinson have since conducted his additional study of Wikipedia articles, finding a strong correlation between the number of edits a Wikipedia article receives and the accuracy and writing quality of the article (Huberman & Wilkinson, 2007). Put another way, the more contributors a Wikipedia article receives, the higher its accuracy and the better the writing. This is not entirely intuitive—certainly, anonymous visitors can and do engage in "graffiti" attempts on Wikipedia pages, and motivated partisans attempt to distort pages to favor their point of view. The site has developed both a set of community norms and practical computer code that lead contributions to have a net-positive effect.

Several authors, including Zittrain (2008), Lih (2009), and Reagle (2010), discuss the details of why Wikipedia manages to succeed. One key attribute is the meager starting set of site rules—(a) articles should display a neutral point of view (NPOV), (b) no copyright infringement, and (c) ignore any rules if they get in the way of building a great encyclopedia—and reliance on the Wikipedia community to mindfully work them out, developing additional rules and protocols as needed.[5] Additional rules have been included over time to manage controversies and improve site quality, but these principles remain at its core. The wiki software code and abundant server space are necessary conditions for this organizing structure. The code lets any community member or passerby offer positive contributions, small or large, while saving past versions for easy review. Graffiti attempts

or biased contributions to an article can thus be removed from the page with a simple click of the "revert to past draft" button. "Bias" and "neutrality" are, of course, terms of art rather than an exact science, but the second attribute helps the community to approximate neutrality rather effectively (see Viegas, Wattenberg, Kriss, & Van Ham, 2007, for further discussion).

A second attribute is the inclusion of a "discussion page" alongside every main Wikipedia page. This is a space for Wikipedians to explain and justify their changes, discuss article quality, and engage in deliberation and disagreement over controversial topics without cluttering the main page. Major edits made without explanation and justification are likely to be reverted, providing an incentive for thoughtful, deliberative engagement. Given the participation of hobbyist communities, many heated "flame war" exchanges occur over topics that are obscure to the mainstream, but passionately debated within a community-of-interest. This is an example of what Lawrence Lessig (1999) terms "code-based governance." Within cyberspace, many of the decisions about how people can and should interact are determined not through government regulation, but by the development of supportive code. Indeed, the original wiki platform did not feature such pages, and after substantial discussion and debate over Wikipedia's listserv, community member Clifford Adams customized the software to create these pages (Lih, 2009, p. 65–66). One challenge for scholars interested in studying the Web's impact on society (a challenge reflected in hypothesis 5, below) is that new code is constantly being developed, and the seemingly impossible dilemmas of 2002 are rendered easily solvable by the new software architecture of 2009. Without discussion pages, Wikipedia would face steep challenges in supporting the NPOV norm. Rather than developing complex organizational bylaws and chains of command, Wikipedia and other online spaces incorporate new code-based solutions that support community norms by making positive contributions easier and negative contributions harder.

The third attribute of Wikipedia's success is the core of initial editors—what I will refer to subsequently as an actively engaged set of "lead adopters." Wikipedia needed this initial group of committed, substantively knowledgeable, and technically skilled contributors because the value of the site is almost entirely derived from its network externalities. Consider the value of Wikipedia to the fifth contributor to visit the site compared to its value to the 5,000,000th contributor. Early on, the site is error-prone, full of topical holes, and of questionable quality. Later, it benefits from a phenomenon first described by IInternet ethnographer Eric Raymond (2001) when discussing the success of the open source software movement: "Given enough eyeballs, all bugs are shallow" (p. 30). Raymond had found that open source software is successful in direct proportion to the size of its community, because a software bug that seems tremendously difficult to one person is likely to be a simple fix for someone else (see Reed, 1999, for further discussion of network effects).

Jimmy Wales explains the success of Wikipedia in similar terms:

> The technology required for Wikipedia is essentially rather simple. You need a database, you need a Web server, you need a Web browser, and you need the wiki editing concept. While the wiki concepts was invented in 1995 by Ward Cunningham, Wikipedia didn't start until 2001. So all of the technology, including the idea of a wiki, which is a Web site that anyone can edit, has existed since 1995. . . . *The answer is, Wikipedia isn't a technological innovation at all; it's a social innovation*. What we figured out between 1995 and 2001 was not new technology. We had the Web already, but we discovered the basic idea of how to organize a community. (Lih, 2009, p. xvi, emphasis added)

This notion of Wikipedia as a community, rather than a technological innovation, is of central importance for generating hypotheses about "open source politics" more generally. As the site has grown, it has incorporated additional rules, and it has empowered a layer of "superusers" with additional editing privileges as a reward for their positive contributions and as a means of engaging in distributed community

management. Wikipedia as a product relies on Wikipedia as a community, engaging on commons-based peer production as a process. The progression through which the community developed is a puzzle that has attracted far less attention than the product or the process, however.

It bears noting that, as suggested by Lev-On and Hardin (2008), the great majority of Wikipedia visitors *do* in fact free ride. Wikipedia globally has about 75,000 "active" members. These are registered users who provide five or more edits to the site per month. About 10 percent of these are "very active" Wikipedians, contributing 100 or more edits per month (Reagle, 2010, p. 6). Given the site's overwhelming popularity, with 8 percent of *all* Internet users visiting daily,[6] we can extrapolate that for every active content-producer, there are tends of thousands who free ride on the public good. Most users of Wikipedia do not take part in the editing or article-writing process, despite the tremendously low barriers to entry. So free ridership does indeed occur on Wikipedia, but it is not the problem that we would be led to expect. No one would likely say that the central issue for Wikipedia is that it is *under*provided.

The key transition in the online space is that, when the costs of participation in collective action approach zero, we face a condition of *abundance* rather than one of scarcity. People have limited time and limited money, but they have virtually unlimited opinions. What we see on Wikipedia is essentially a multifaceted version of what Olson termed a "privileged group."[7] When the resource in question is not money or time, but rather specialized information, we find that there are plenty of people who are "wealthy" in some form or another. Put another way, most everyone has a hobby. Hobbyists have always happily incurred the "costs" of discussing their topic of interest, often in excruciating detail. When they do so on Wikipedia, they provide exactly as much of the public good (information about video games, the history of knitting, etc.) as they themselves want, and this provides more than enough for inquiring minds.

This is not to say that mass collaboration, collective action, and the provision of online public goods is seamless and assured. Rather, it is to say that the shift from slower, costlier information regimes to an instantaneous, abundant online information regime creates a *different* dilemma for social engagement. Specifically, the geography-less, abundant online space creates tremendous challenges in *search*. How are we to identify good, verifiable information from bad? How are motivated partisans or hobbyists to find each other with no central square, and how are onlookers to take advantage of the fruits of these hobbyists' labor? Wikipedia critically benefits from the *network externalities* of all these hobbyist communities gathering in the same, identifiable location. If five sites all competed for the same niche of "online organizational hub" (Lev-On & Hardin, 2008, p. 16), the sum of those parts would be far less than the whole found on Wikipedia. Indeed, initial developer Larry Sanger eventually left Wikipedia and started his own site, Citizendium.org, because he felt there should be a greater role for credentialed experts (Bruns, 2008; Reagle, 2010). In two and a half years, the site has built a small community of 800 contributors, authoring 10,700 articles in total and attracting a fraction of a percent of Wikipedia's audience.[8] For this reason, I depart from Lupia and Sin (2003) and Bimber et al. (2005). I would suggest that the critical challenge to online collective action is not public–private boundary-crossing or the declining value of selective incentives, but rather solving the search dilemma under conditions of abundance—a challenge that approximates a mass coordination game.

HYPERLINKS, HUBS, AND POWER LAWS: AN ITERATED SOLUTION TO THE SEARCH DILEMMA

Before there was the World Wide Web, there was the hyperlink. Hyperlinks provide the networked structure of the Internet, with clickable links embedded in text that direct a reader from one page of text to another. A solitary Web page with no inbound or outbound hyperlinks lies, in a very real sense, at the periphery of the World Wide Web. Though such a page is accessible through direct input of its uniform

resource locator (URL: the text-based "address" appearing after http:// in the address line of a Web page), one would be unlikely to stumble upon it through everyday surfing.

The hyperlink calls to attention two dimensions of the Internet's novel search puzzle. First is the anti-geographic nature of the medium itself. Search in the offline world is aided by landscape-imposed scarcity. Towns and cities have physical centers and peripheries, and this translates directly into the price system of the real estate market. There is a cost imposed by being out-of-the-way, either for residencies (commute) or commercial zones (foot traffic and shopping districts). Thus restaurants tend to be grouped together, one can generally expect to find a pawn shop in close proximity to a race track, and proximity to desirable locations translates into higher rents. On the Internet, by contrast, there is no physical landscape to traverse. As one example, consider the hundreds of millions of blogs that have been created and then abandoned. This provides the slightest inconvenience for Google, the company upon whose server farms most of these sites are hosted, and whose search algorithm must handle them, but the realities of increasing bandwidth and transistor capacity relegate this to a minor nuisance at most. From the user's perspective, dead blogs and abandoned Web pages do not litter any landscape, because the Web is composed of hyperlinks and we are never forced to traverse their pages in our daily online pursuits. An abandoned blog goes unhyperlinked, and thus floats to the periphery of Web "space." The lack of geography on the Web is a substantial component of the condition of information abundance found online. There is no such thing as "location, location, location."

The second dimension is the challenge for like-minded hobbyists of finding *each other*. Internet communication is instantaneous, but also asynchronous. One can post a message to a discussion board or send an e-mail alert and it will be immediately viewable, but as opposed to a phone or face-to-face conversation, replies do not necessarily come in real time. Lacking town centers, where are hobbyists, partisans, or other communities-of-interest to gather? With no town center, what good is a self-publishing soapbox, anyway? This is closely related with the problem of identifying verifiable information on the Web. In essence, the Internet lowers the communication costs for *all* types of publication and online group interaction. Scarcity provides some baseline assurance that a group or information source is reliable; the very act of publication or gathering indicates an ability to surpass some minimal cost threshold. Under the condition of abundance, how are we to tell reliable information from speculation? How are we to find other like-minded participants when there literally is no "there" there?

Hyperlinks provide the kernel of the solution, with Google's PageRank algorithm acting as pioneer. Prior to PageRank, Internet search was tremendously problematic. The two standard solutions were to provide a top-down directory of all Web pages or offer a search mechanism based on the appearance of keywords on a Web page. The problem with directories was twofold. First, the scale and rapid growth of the Web meant that no directory could manage to be comprehensive. Second, directories are built around meta-level organizing assumptions about the categories a user will wish to search through. Thus AOL.com, for instance, could provide a list of topical headings such as "sports," "news," and "entertainment" and then further divide the categories into fine-grained subheadings. But a user interested in new banjo strings and information on an upcoming jamboree would have little idea where to begin. Keyword-based search could help with this, organizing results based on the combination of "banjo strings" and "jamboree," but separating new information from old becomes problematic, and such keyword searches are easily gamed. Google's ingenious solution was to include hyperlink data in the presentation of search results. Pages with numerous hyperlinks, particularly from other sites that are highly linked, appear at the top of the results page. Thus Google lets Web users "vote with their feet," in a sense—indicating the quality of an information source based on the number of Web users who have chosen to link to it. The simple inclusion of this network data in its search results is what led Google to rise from a tiny startup, three-person operation to the largest company in the online space (Vise & Malseed, 2005).

Physicist Albert Lazlo Barabasi offered an important early treatment of these link patterns on the Web in a 1999 article in *Nature* magazine. As he would later describe in his public-audience book, *Linked*, Barabasi was interested in the distribution of links among Web pages. His early assumption had been that link distribution would approximate a normal curve, indicating that the Web could be understood mathematically using the standard assumptions of random graph theory. Instead, Barabasi found that link patterns followed a heavily skewed distribution approximating a power law or Pareto distribution. Vilfredo Pareto initially observed these distributions in his study of wealth disparity in European societies, leading them to often be termed "rich get richer" or "80/20" distributions, since he found that 80 percent of a society's wealth was held by the top 20 percent, and that the greater the level of income, the more stark the disparity. Power laws are based on a decaying function in which the Nth-largest node is 1/Nth the size of the largest node (Barabasi, 2003). Shirky and Hindman, Tsioutsiouliklis, and Johnson produced separate studies in 2003 demonstrating that the blogosphere in particular displays power law tendencies in its hyperlink distribution, leading to the emergence of an "A-list" or elite status among early political bloggers. Karpf has likewise found the same pattern evident in the right- and left-wing blogospheres, noting that each blog community displays *its own* power law distribution (Karpf, 2008). Though there has been some debate as to whether these link patterns are a power law or some other heavily skewed distribution (Drezner & Farrell, 2008), what is of particular interest here is the *mechanism* that Barabasi tells us produces power law distributions.

Barabasi demonstrates in his article that power law distributions emerge in a network simulation when two simple conditions are present: (a) growth and (b) preferential attachment. Simply put, if a network is growing and new links between nodes are determined based upon the preferences of their observable neighbors, then a set of "hubs" will develop over time, as the link-rich are more likely to gain additional links, further increasing link disparity over time and, critically, developing a power law distribution. Growth plus preferential attachment leads to the emergence of power-law hubs. In so doing, this type of hub formation also serves as an iterated solution to the mass coordination problem found online.

Let's say you are interested in discussing left-wing politics. Living in a conservative rural town, you would like to turn online in order to find other people with similar interests. Where do you go? Where are they? The previously mentioned lack of geography provides a dilemma. You have no strong preference regarding the location of the conversation, and neither do the other members of your nascent community-of-interest. Your interest is in finding the same "place" online (and, later, in the place providing supportive environment for healthy, spam- and "troll-"free discussion and possibly tools for further collaboration). This is a classical example of a coordination game, in which actors have neutral preference rankings among options, but wish to arrive at the same solution as one another. In a single-iteration coordination game, this can be solved through sequential action: The first actor makes an arbitrary decision, and all others follow suit. If actors move simultaneously, or without knowledge of each other's actions, the problem becomes far more challenging. But in an iterated coordination game, preferential attachment emerges as a viable and simple solution. In particular, a Google search will reveal the most popular spaces where like-minded people are *already* meeting. Rather than selecting an online forum, blog, wiki, etc. at random and hoping that a community-of-interest will show up, each additional latent community member can choose to rely on the actions of those who came before him or her.

Preferential attachment leads directly to the emergence of power-law hubs, and a general principle for Web science practitioners: Large hub spaces online are *different* than small spaces. The topology of the Web, as it has grown over time, is ruled by power-law hubs such as eBay, Wikipedia, DailyKos, YouTube, MoveOn, and Facebook. Each of these "Web 2.0" spaces offers value to its users in direct proportion to the network effects provided by large crowds of similar users. Online hub spaces move through

identifiable phases of institutional development as they diffuse through the user population and face challenges related to scale and changing demographics and interests of different user classes.

The study of power law distributions in online politics is mostly attributable to Matthew Hindman's work, particularly his 2008 book *The Myth of Digital Democracy*. Therein Hindman argues that the emergence of power laws in online traffic creates a "Googlearchy," or Google-imposed hierarchy, leading to heavy elite stratification and limiting the transformative potential of the medium. Working with traffic data supplied by *Hitwise*.com, Hindman argues that the barriers-to-entry online are not substantially lowered by the new media environment. Though the costs of self-publication have been dramatically reduced, those costs have been offset by the new costs of building a massive online audience. To Hindman, these power laws are a problem; he argues that they represent the re-emergence of elite politics that in turn limits the transformative potential of online communication.

Without disputing Hindman's empirics, the case of Wikipedia suggests that we should be circumspect about his interpretation. The Internet's power-law topology means there can only be one hub site occupying Wikipedia's niche. But Wikipedia's users are not attempting to build Wikipedias of their very own. They are, instead, looking for a coordination point where they can access "the sum of all human knowledge." The path to power-law hub status is a developmental process, and it yields a set of institutions that are substantially more open and participatory than those characterizing the previous information regime. The utility of power-law hubs in solving the mass coordination problem has been largely ignored in the research literature thus far.

INSTITUTIONAL DEVELOPMENT OF HUB COMMUNITIES: A FIVE-STAGE ADOPTER CLASS MODEL

Wikipedia benefits from the power-law topology of the Internet, developing a large community of participants, active and passive, and benefiting from the substantial network externalities that they provide. The rise from nascent startup to power-law hub did not occur in a smooth progression, though. Wikipedia was able to succeed because its leadership skillfully and artfully moved it through a predictable series of development challenges that occurred as the community grew and changed. All such Internet-mediated community spaces move through the same diffusion process as virtually any other new product or innovation: (a) beginning with a tiny group of lead adopters who co-create the good, (b) expanding to a larger early adopter class, which is highly motivated but less technically skilled, (c) launching into the much larger early majority class, whose motivation and skill level is more varied and whose size pressures the system to adapt, (d) adopting protections against spammers and malicious attacks as the site attracts the late majority class and becomes recognized as "valuable online real estate," and (e) dealing with challenges to institutional power structures as growth slows at the laggard phase and questions regarding voice and equality rise to the fore. These stages are of particular interest because they accord both with Wikipedia's experience and with the longstanding literature on diffusion of innovations (Rogers, 2003). If Hindman and others are correct about the stability of power-law hub sites online, then there can only be a small number of these online communities-of-interest, and their development pattern is itself an important topic for investigation.

S-CURVES AND ADOPTER CLASSES: A BRIEF OVERVIEW OF THE DIFFUSION LITERATURE

The definitive text regarding diffusion research is *Diffusion of Innovations* by Everett Rogers. First published in 1962, the book is now in its fifth edition and has been cited over 19,000 times[9]—a testament to Rogers's longstanding impact on the field. Rogers notes that ideas, farm products, viruses, and a whole range of other innovations fit a standard "S-curve" as they diffuse through a community over time.

FIGURE 1. Standard diffusion curve[13] (color figure available online).

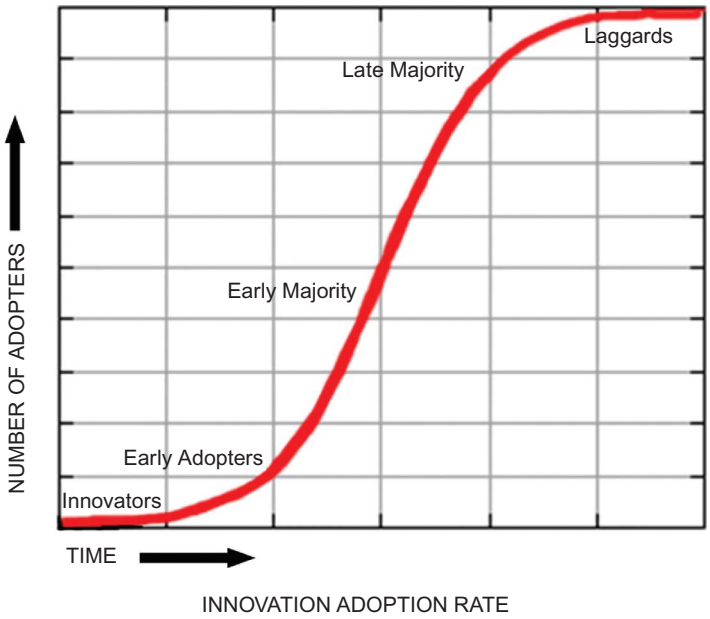

Figure 1 offers a graphical representation of the S-curve, along with the five traditional adopter classes. Eric Von Hippel (2005) re-labels the "innovators" as "lead adopters" in his book, *Democratizing Innovation*. He notes in that work that the first tiny group of adopters often helps to co-create the good, repurposing it and providing feedback to the original firms or labs who are releasing the new product. This is particularly true in the computer industry, with beta-testers providing feedback to proprietary software companies and open-source programmers actively participating in the software development process. Following Von Hippel, I use the term "lead adopters" rather than "innovators" here. Note the relative size of the five adopter classes, with lead adopters being the smallest group, the early and late majorities making up the bulk of the population, and early adopters and laggards representing 13.5 percent and 16 percent of the population apiece, respectively. This is based on an assumption that time-of-adoption follows a normal curve, with the early and late majorities covering one standard deviation from the mean, early adopters representing the second standard deviation to the left of the mean, lead adopters representing 2+ standard deviations to the left and laggards representing all adoptions occurring more than one standard deviation to the right (Rogers, 2003, p. 281).

One of the most robust findings from the diffusion literature is that these adopter classes are demographically distinct from one another. Survey research has routinely found that younger, wealthier, better educated, and more "cosmopolitan" members of society have a stronger taste for innovation than their neighbors (Rogers, 2003, pp. 272–282). Lead adopters and early adopters tend to have peer networks that span wide geographies, exposing them to new ideas and innovations long before their neighbors do. Thomas Valente (1995) further advances this notion of separate adopter classes in *Network Models of the Diffusion of Innovation*. Valente unites the longstanding diffusion research tradition with the emerging field of social network analysis, treating actors in a community as nodes in a network with varying adoption thresholds. He goes on to identify three critical mass points: one at the shift from early adopters to early majority, a second at the pivot point between early and late majority, and the third at the shift from late majority to laggards.

This approach is particularly valuable because it suggests that not only are there differences between adopter classes, but there are also temporal differences between the various phases of adoption.

It is worth noting at this point a methodological difficulty in the diffusion and networks literatures. As Wasserman and Faust (1994) note in their text, *Social Network Analysis: Methods and Applications*, population definition is a crucial and troubling issue. For early diffusion researchers studying farm implements, the population under study would be farmers in an identifiable community. For later research on drug development, the population would be medical doctors with a shared specialty and overlapping memberships in the American Medical Association. What is the population of potential Wikipedians, though? What about the population of potential Dean campaign participants, or Tea Party activists, or MoveOn members? Boundary definition can only be determined in retrospect for these groups, rendering social network analysis useful for theoretical exercises, but presenting substantial data hurdles for more quantitative work. For this reason, I use the diffusion and social networks literatures as a starting point for my descriptive model of institutional development in online communities-of-interest, but do not develop the model as a social network study per se.

INSTITUTIONAL DEVELOPMENT CHALLENGES PRESENT AT EACH ADOPTION STAGE

The important lesson from the diffusion of innovation literature is that the fifth Wikipedian is substantively *different* from the 5,000,000th Wikipedian. They have different backgrounds, different interests in the site, and different needs of the site architecture. The fifth Wikipedian is co-creating the online space. She is likely involved in writing software code or is particularly devoted to the creation of an open encyclopedia. The five-millionth Wikipedian is visiting an established online space, looking up information written by others, and eventually finding enough value in the space to add a few edits of his own. Effective launch of one of these communities must move through five distinct phases: (a) initial launch, (b) reaching critical mass, (c) managing the influx of a mass public, (d) defending norms against newcomers, and (e) institutionalizing authority. Each stage is discussed below:

Stage 1: Initial Launch

Recall again Jimmy Wales's suggestion that the technology behind Wikipedia was both simple and available for years prior to the launch of the site. The success of Wikipedia was a story of community-building. If Wales and Sanger had announced Wikipedia with an aggressive television and newspaper advertising campaign, the site would have been an almost guaranteed failure. The mass audience would have visited an empty vessel populated by a few anonymous (and likely erroneous) entries, turned around, and never come back. But the initial Nupedia list gave them a small set of highly motivated participants who could choose to contribute to the site because they individually found it a fascinating and worthwhile project. Their "adoption threshold" in the language of Valente, was tremendously low. The site also had the early blessing of Richard Stallman, founder of the Free Software Foundation and legend within the open source software community, and received an influx of tech-savvy participants through early discussion on the Slashdot.org discussion forum, described by Lih as "a salon for the technical elite and a grand senate of the computing community" (Lih, 2009, p. 67).

The attention of this lead adopter community is itself a scarce resource: they are, as a whole, well educated, urbane technology and academic professionals with time for a few interesting side projects and a dense network of social ties. Benkler and Weber both note that the personal incentive for these individuals lies in a combination of reputation-building incentives, socio-psychological incentives, and hedonic personal gratification at solving interesting puzzles (Benkler, 2006; Weber, 2004). Any online community-of-interest must attract a sizeable number of these lead-adopting co-creators, and

that in turn means providing them with the freedom to make changes and provide input to the system. Internet communication may exist in an environment of information abundance, but the interest of these elites is a scarce and valuable resource, best attracted through technology conferences, highly technical listserv discussion groups, and other traditional networking events that feature high barriers-to-entry. Though the identity of the lead-adopter community will vary from online community to online community (the lead adopters who populated the early political blogosphere were not the same people who populated early Wikipedia), they are invariably drawn from existing networks of influence—the underdefined "policy networks" discussed in the policy agenda literature, for instance (Kingdon, 1984).

Stage 2: Reaching Critical Mass

"User-generated content," like Web 2.0, is an Internet buzzword coming out of the marketing world that has taken on substantial meaning. Web 2.0 can be roughly defined as people finding each other online, whereas Web 1.0 consisted of people finding information online ("the information superhighway"). User-generated content refers to comments, information, conversation, or multimedia content that come not from top-down management, but from bottom-up, voluntary production. Several of the major online spaces (circa 2010) serve to aggregate and sort such content, including Facebook (publicly articulated social network information), YouTube (video sharing), Flickr (photo sharing), the large community blogs, and of course Wikipedia. To the extent that Internet-mediated political associations rely on channeling the participation and interaction of their communities-of-interest, they likewise fit into this category. Critical mass refers to the point at which a site is receiving enough user-generated content that the network externalities produced exceed the interest threshold for the mass of less-motivated Web surfers. Put another way, at some point Wikipedia has enough content to maintain the interest of people who do not self-identify as "techie" or "encyclopedia junkie." The addition of this larger swath of the public massively expands the community. As the nascent version of any of these sites expands into this early adopter phase, it must settle a series of technical and normative questions regarding how to handle growth and community contribution.

In Wikipedia's case, this included some complicated server load issues (Lih, 2009, pp. 77–79) in 2004, as the number of total English-language articles surpassed 100,000 (about the size of the *Encylopedia Brittanica*) and increased traffic grew to the point where the site would often crash. The involvement of technical elites was critical to solving these problems, and all growing online communities must either attract the sustained interest of the open source community or maintain a large budget for proprietary software solutions to this aspect of scaling. Lih records that, in the same time period, "because the community was growing so quickly, the process of forming consensus by e-mail did not scale" (Lih, 2009, p. 95). The consensus and co-creation practices that were necessary to attract and keep the lead adopter community had to be modified in order to allow for the early adopters, who by and large displayed a keen interest in the system, but were less technically experienced and lacked deep existing network ties with one another. Wikipedia responded by creating a distributed moderation system of superuser "administrators" and by moving mailing list–based discussion to a separate section of the wiki dubbed the "village pump." As Wikipedia attracted enough user-generated content to become self-sustaining, then, the system had to adopt new code-based solutions to the surge of traffic.

Stage 3: Managing the Influx of a Mass Public

As the site reaches Valente's (1995) first critical mass point, it must deal both with a tremendous surge in traffic/participation and also adapt to a mass public that does not share the particular interests of the lead and early adopters. While lead adopters are contacted through existing social/professional network ties, and early adopters are contacted through niche media

outlets (coverage in *Wired* magazine being particularly coveted by many social media ventures at this stage), the shift to early majority is often accompanied by coverage in traditional media venues. Wikipedia had attracted a few brief mentions in the mainstream media during its first few years, but its breakthrough moment occurred during a well-publicized controversy in December 2005. John Seigenthaler, former editor of *The Tennessean* newspaper, noticed some incorrect and libelous information posted in his Wikipedia entry. Seigenthaler contacted the editors, who immediately changed it and apologized, but Seigenthaler went on to write a scathing op-ed for *USA Today* on Wikipedia's unreliability regardless. The op-ed produced ripple effects, with other television and newspaper outlets writing stories about the story (Seelye, 2005). For millions of Americans, this coverage was their first introduction to the site's existence, and the negative news served as free site publicity that significantly increased traffic and content-creation.

In the history of Wikipedia, this is referred to as "the Seigenthaler effect." Figure 2 demonstrates the growth in Wikipedia page views pre- and post-Seigenthaler op-ed. The upward trend in views continued unabated, as Wikipedia grew to its present-day status as the sixth most-visited Web site in the world. This sustained growth would not be possible prior to the normative and technical problem-solving occurring in stage 2—the site would lack a vibrant community and also lack the capacity to deal with the sudden influx of users. As-is, the arrival of the early majority signaled a change in the character of the site, as the culture of "ignore any rules that get in the way" had to stand up to the rush of onlookers less sure of their co-creating skills and more interested in a simple set of guidelines for what can and cannot be done. It is generally during this third stage that many of the lead adopters, faced with changing community norms and an increasingly noisy environment, depart for some new project or create their own sublist, complaining about how the community has been degraded by the onrushing newcomers (Shirky, 2008).

FIGURE 2. Growth of Wikipedia page views pre- and post-Seigenthaler op-ed[14] (color figure available online).

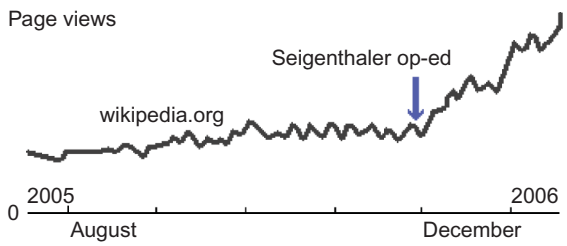

Stage 4: Defending Norms against Newcomers

As the online community passes Valente's second inflection point, growth is at its highest rate and the network externalities have rendered the space a clear power-law hub. At this point, the site becomes known as "valuable online real estate." A new wave of challenges comes with such a distinction, as malicious users attempt to subvert the network for their own gain. Wikipedia has remained surprisingly robust against these challenges—a credit both to the technical solutions it has created and the participatory community it has enabled. But two examples of this challenge demonstrate the general point. On July 31, 2006, political humorist Stephen Colbert featured Wikipedia in a segment of his television show, *The Colbert Report*. Describing Wikipedia as a space where, "any user can change any entry, and if enough users agree with them, it becomes true," Colbert told his viewers to go onto Wikipedia and edit the article on elephants to say: "Elephant population in Africa has tripled over the past six months." The flood of user-edits forced site administrators to temporarily lock the page. In a less congenial spirit, companies and political aides have gotten into the habit of anonymously grooming their entries. Zittrain (2008) elaborates the tension admirably: "If the Wikipedia entry on Wal-Mart is one of the first hits in a search for the store, it will be important to Wal-Mart to

make sure the entry is fair—or even more than fair . . .” (p. 139). Likewise, August 2006 saw the launch of MyWikiBiz, a company aimed at creating and editing Wikipedia entries on a for-fee basis. Jimmy Wales responded by blocking the company’s user account and banning its IP address, and this led to a lengthy community discussion about how to deal with such new ventures (Zittrain, 2008, p. 140).

The “valuable real estate” issue has important implications for the growth of online communities in areas that have already been identified as valuable. When the *Los Angeles Times* attempted to embrace the wiki editing concept through the launch of “wikitorials,” the site was almost immediately overrun by porn advertisements and was quickly shut down. Clay Shirky (2008) writes, “In the absence of a functioning community, a wiki will suffer from the Tragedy of the Commons, as the Wikitorial did, as individuals use it as an attention-getting platform, and there is no community to defend it” (p. 137). Karpf (2009b) identifies a similar trend impeding nascent conservative online political communities in their efforts to build parallel infrastructure to the progressive “Netroots.”

Stage 5: Institutionalizing Authority

Throughout the first four growth phases, we see a continuous fraying of the principles of openness and co-creation that mark the earliest stages of a participatory community. As sites enter the laggard phase (which I will again note, can only be methodologically defined with rigor retrospectively), the slowdown in site growth raises inevitable questions of power and authority among the now-stabilizing community. Within Wikipedia, one such controversy occurred when longtime site administrator “Essjay” was revealed to have falsified his credentials. Although Wikipedia is open to editing from anyone, Essjay had claimed on his personal page that he held various graduate degrees and a professorship in theology. He had made reference to this educational background when arguing on various “talk” pages over the years. In 2007, after Jimmy Wales contacted him about joining a for-profit venture, it turned out that Essjay was a 24-year-old editor with no graduate degrees. This led to a long community discussion regarding the validity of his edits, the issues of identity-management in the online space, and the proper role of expertise in Wikipedia (see Zittrain, 2008, p. 141; Lih, 2009, pp.194–200 for further discussion).

As growth slows in this final phase, when most potential community members have joined the site and the remainder of the online population is mostly non-adopters with a few laggard adopters still present, the disparity between hubs and niches comes into stark contrast. While the periods of rapid growth provide a sense that the entire world is changing, the final phase raises questions about who controls the fruits of all this volunteer labor. These changes have been somewhat muted in Wikipedia because the site is a nonprofit, nonpolitical venture. But in other communities-of-interest, particularly ones where a company or political leadership is seen to profit from the voluntary output, the challenges to institutionalized authority can be particularly problematic. The differences of scale that have developed become differences-in-kind, with Larry Sanger’s attempt to start his own equivalent to Wikipedia, Citizendium.org, being an instructive case. As Internet publisher Tim O’Reilly has put it, “If there weren’t a network effect driving Wikipedia, [Google’s] Knol and Citizendium would be succeeding.”[10] The powerful network effects that define these online spaces also prevent alternative ventures from successfully growing to scale. If you don’t like Wikipedia, DailyKos, or Facebook, you are free to start your own, but that in itself is problematic.

If the power-law topography creates these differences-in-scale among the sites that allow for novel solutions to the collective action problem, then we must wonder about the conditions under which a power-law hub can fall or be replaced. The next section will discuss how each of the five institutional development stages listed above produces a challenge that can lead to the failure or replacement of a network-enhanced

good, leading to a set of hypotheses about the parameters within which potential network-enhanced goods can be developed.

STUMBLING ALONG THE PATH TO POWER-LAW HUB STATUS: THE LIMITS OF OPEN SOURCE POLITICS

Phase 1

The first challenge for developing a hub space lies in attracting a devoted set of lead adopters. This problem can come in at least two forms, depending on the availability of a pre-existing power-law hub. In the case of Wikipedia, for instance, the first wave of adopters came from the Nupedia list and from the Slashdot community. Likewise, the Howard Dean campaign featured the support of dozens of leaders in the field of social technology who were attracted by the new opportunity to apply the principles of open source to a political campaign and see what happened. Those are existing networks whose attention is a scarce commodity. Attempts at replicating these successes must find some other reason why an existing technological or topical elite network would choose to engage in peer production through that particular venue.

This point seems lost upon the hundreds of organizations and companies who have decided to enter the Web 2.0 age by launching their own social networking sites, for instance. A useful indicator is the existence of a McDonald's social networking site. Millions of Americans eat at McDonald's, but how many of them wish to self-identify as members of the "McDonald's community"? Pushed forward by a consulting industry that has found lucrative contracts in supporting the growth of social media, the very real public goods produced by online communities-of-interest can be easily obscured if we look at the social media industry as a whole. Without a colonizing set of devoted, skilled volunteer participants, the best technology in the world will fail to deliver the valuable network externalities that make these spaces worth regularly visiting. In a similar vein, I argued in a recent *Journal of Information Technology & Politics* article, "Macaca Moments Reconsidered," that the primary impact of new media on politics is only identifiable through the study of large-scale communities-of-interest, rather than the isolated study of specific new media tools such as YouTube (Karpf, 2010). The impact of the Internet on political associations and politics in general comes not through lowered communication costs alone, but through the communities-of-interest that these lowered costs enable. Open source as a process is only as effective as the attached community makes it. The first step in building such a community lies in attracting a set of active co-creators, and these co-creators are themselves a scarce commodity. This leads to the first testable hypothesis, also termed the "Field of Dreams Fallacy":

> H1 (The "field of dreams" fallacy): Successful peer production requires the initial engagement of a lead adopter community, predominantly found within a pre-existing network.

If H1 is correct, it follows that optimism about "open source politics" must be tempered by a recognition of existing power relationships. It simply is not the case that "if you build it, they will come." Rather, existing networks of elite actors must be courted and brought aboard. Whether these are members of the "technorati" or other social, political, or media elites, they represent a pre-existing set of "haves" whose participation is a necessary condition for the success of even such open, egalitarian architectures as the one found on Wikipedia.

Phase 2

The move from lead adopters to the larger set of early adopters represents a distinct bundle of challenges. Lead adopters are a valuable commodity, but they also have many interests that are quite different from the rest of the population. Reaching critical mass requires that a site not only solve a series of technical and normative challenges; it also requires the new community to exist in an area which is attractive to a substantial issue public. Shirky writes about a variant on this hurdle in his 1999 essay, "The

Interest Horizons and the Limits of Software Love." Responding to Eric Raymond's then-recent summary of open source, that "every good work of software starts by scratching a developer's personal itch . . . given enough eyeballs, all bugs are shallow," Shirky notes, "What if you have a problem that *doesn't* scratch some core developers personal itch?" (Shirky, 1999). Within the restricted universe of software development projects, some ideas will be more exciting and motivating than others. The least exciting ideas may still have a commercially viable market, but they are unlikely to attract a large enough community of motivated developers to be appropriate for commons-based peer production.

This critique holds for the formation of online communities-of-interest as well—which is not surprising, given that Wikipedia and other such communities took inspiration from the open source software movement. The lowered transaction costs of the Internet help to reveal the full demand curve for public participation, but part of what that means is that topics or areas that simply are not particularly attractive or interesting to any existent or nascent issue public will fail to reach critical mass. Revelation of a complete demand curve does not mean all issues will be equally in demand.

The first generation of social scientists to study the Internet was optimistic that, thanks to the falling costs of online engagement, we would see the rise of mass deliberative spaces, "online public squares," and other venues for enhanced democratic participation. Many such sites have been launched with enthusiasm, only to fail to reach critical mass. There are several potential explanations for such failure, but one of them is that public interest in lengthy deliberative processes simply is not as high as social scientists would ideally like. (See Schudson, 1999, for a similar historical discussion.) One limit of peer production that will hamper communities-of-interest is the inability to attract a large enough community to pass the critical mass point where the user-generated content itself gives people a reason to regularly return to the online space. This leads to a second hypothesis, termed the "interest horizons" thesis:

> H2 (The "interest horizons" thesis): Commons-based peer production will extend beyond the lead adopter community only to the degree that it appeals to the interests of a broader community.

The interest horizons thesis may sound nearly teleological in nature—"peer production communities will only form where they are interested in doing so"—but it offers an important corrective to the practical limitations on open source politics. The Obama campaign, for instance, featured many "open source"–type activities. Will the relative successes of a site like MyBarackObama.com be replicable for a local county council candidate in 2010, however? H2 would suggest not, for the simple reason that presidential campaigns attract much greater public attention than local races. Presidential campaigns—Obama's in particular—are much higher on the *demand curve* of the American public than any other electoral campaigns, and thus there are "open source" strategies that can only be successfully applied in such settings. The promulgation of "best practices" white papers and anecdotal accounts of success largely ignores the implications of the interest horizons thesis.

Phase 3

Often launched by some event that exposes the hub space to the population through the mass media, the third phase is where substantial scaling and network effects begin to take hold. An important related challenge at this juncture is the availability of a distributed reputation system capable of managing this scaling process. As discussed by Benkler (2006), Bruns (2008), Resnick, Zeckhauser, Swanson, and Lockwood (2006), Karpf (2011), and others, online reputation systems are a necessary component of all hub spaces within the power-law topography of the Internet.

A "benevolent dictator" such as Jimmy Wales can play a guiding role in the first two phases of growth, but in phase three, communities of interest quickly learn that "Jimmy doesn't scale" (Lih, 2009, p. 179). Slashdot's "mojo" system and eBay's "feedback forum" are the two

best-known examples, but Google's PageRank algorithm has similar functions, drawing upon a large set of distributed reputation assessments, then applying some form of algorithm that rewards good content or contributions while sanctioning bad content or contributions. Yochai Benkler (2006) notes in *The Wealth of Networks* that an effective reputation system is a necessary condition of large-scale peer production. He goes on to suggest that the components of peer-produced systems can be broken down into smaller components ("modularity") and that these components themselves can then be reduced to tasks that require little time and effort ("granularity") (Benkler, 2006, p. 100). Benkler illustrates these points by drawing upon the set of existing online reputation systems, but in so doing he overlooks an important caveat: Some types of mass collaboration are much more easily reduced to small actions taken in front of a computer monitor than others.

This represents a substantial limitation to the Internet's impact on political associations. Wikipedia, DailyKos, MoveOn, and other large-scale communities-of-interest are capable of overwhelming growth with low overhead costs because they are asking their community to engage in distributed tasks that can occur effectively in front of a computer screen. One challenge that MoveOn, Democracy for America, and similar organizations have faced when they attempt to use "online tools for offline action" is that the slight increase in transaction costs—asking people to rate meetings after they return home to them, for instance—is accompanied by a steep drop-off in participation.

Karpf (2011) argues that these limits are changing thanks to the diffusion of the mobile Web (i.e., Internet-through-iPhone), but it is still too early tell whether that hypothesis will be supported by the eventual data. For our purposes here, it bears noting that the impact of the Internet on offline collaborations is slim when compared with its impact on online collaboration. Potential power-law hubs can only radically scale up if they adopt a system to manage the influx of participation. Such systems of reputation and recommendation are not equally applicable to all forms of collaboration and engagement, and where they cannot yet be built, commons-based peer production will fail to displace traditional modes of association and production.

> H3 ("Political strategy is not like computer code"): Commons-based peer production is limited to those tasks or activities that are amenable to online reputation systems.
>
> H4 ("Location-based wave" thesis): The spread of the "Mobile Web" will expand the range of tasks or activities amenable to online reputation systems.

Phase 4

By the fourth phase, a site has managed to attract mass attention and benefits from substantial network effects. What is to stop it from continuing in this regard? The brief history of social network sites (SNS) offers a useful illustration. Friendster.com was the first SNS to pass critical mass and attract large-scale participation. danah boyd chronicles the demise of Friendster, eclipsed by Myspace.com because MySpace offered a more permissive culture, inviting bands to start their own pages and letting users create fake profiles for schools, organizations, and celebrities. Friendster had a network externality-advantage, because more people were initially on its site, but low online transaction costs meant that people could add a MySpace account in minutes, and with greater freedom on MySpace, they eventually switched en masse. boyd attributes the replacement of Friendster by MySpace as an indicator of "Internet culture" (boyd, 2006; boyd & Ellison, 2007).

MySpace indeed gained millions more users than Friendster, as SNS's gained further penetration among the public at large. Matthew Hindman notes that, prior to June 2007, MySpace was stably among the top five Web sites in the United States. In his research into the stability of power laws on the Web, he notes that MySpace precipitously dropped that June because the site "became uncool."[11] In the months leading up to that decline, MySpace had become barraged by spam solicitations, as pornography marketers took note

of its status as "valuable online real estate" and began creating fake accounts. Viruses also became a problem around this time. Critically, Facebook.com replaced MySpace at this time, and it remains the SNS power-law hub today. Facebook included more limiting user registration permissions, and only allowed members of an individual's school-based or geographic network to view his or her profile. Perhaps more importantly, in May 2007, Facebook unveiled a new feature: its open application programming interface (API). The open API allowed outside developers to write new programs, including games and information-sharing tools.

Facebook replaced MySpace as power-law hub, not because of culture, but because the open API gave users *something new to do*. Failure to respond to the pressures of being "valuable online real estate" rendered MySpace vulnerable, and when Facebook gave users new engagement opportunities, MySpace was left as a virtual ghost town, with over 100 million registered users, most of whom were suddenly spending the bulk of their time on another site.[12] (See boyd 2009 for a more detailed argument regarding the socioeconomic effects of this "Myspace flight.").

The lesson we should draw from the history of social network sites is that, although power-law hubs benefit from substantial network effects that render substantial stability in the short run, viewed over a long time horizon, the hubs appear more fragile. The Internet is a fast-changing environment, and lead adopter communities tend to desert an online space once it gets too noisy and crowded, moving on to experiment with the next wave of innovations. Just as Compuserv, AOL, and Geocities were once defining features of online "geography," only to be relegated a few years later to the dustbin of history, the changing nature of the Internet creates room for a host of "disruptive innovations" (see Karpf, 2009a, for further discussion) that can lead to the displacement of existing hub communities. This is a particularly pronounced feature in light of "Moore's Law," which predicted in 1965 that transistor capacity would double approximately every two years (Moore, 1965). Unlike other previous communications regimes, online communication continues to evolve at a tremendous pace, creating constant opportunities to introduce new features to a system that had been cost-impermissible in the previous election cycle. This leads to H5:

> H5: ("Power law fragility under Moore's Law thesis"): The displacement of an existing power-law hub by a new entrant will be facilitated by the development of novel features, enabled by Moore's Law.
>
> H6 ("Punctuated equilibrium thesis"): Viewed over time, individual topic areas online will demonstrate "punctuated equilibrium" characteristics, with long periods of stability, followed by a power-law hub being replaced by a new upstart.

H5 suggests that, as the Internet continues to evolve as a medium, power laws may turn out to be less stable than they at first appeared. New startups invest in capacities of the Web at time T that were not present at time T – 1, and in the lowered transaction cost digital environment, this can lead to the replacement of one hub with another. The requirements of the mass coordination dilemma require that the environment will continue to only have one central hub per area, but those hubs can be replaced in a predictable manner. Likewise, H6 posits that the apparent stability found in the online landscape may bear a strong resemblance to the stability found in the policy subsystems literature (Baumgartner & Jones, 1993). At a single point in time, policy subsystems appear stable and unchanging. Viewed over a 30-year timeframe, they instead reveal short periods of disruption, followed by the establishment of a new systemic hierarchy.

Phase 5

The governance challenges presented in the fifth and final stage are difficult to describe in great detail, particularly because of the data limitations present when applying social network analysis to online communities-of-interest. One cannot say with certainty whether Wikipedia or more political hub spaces have actually entered the laggard phase of adoption, because we do not know at present what percentage of the online population is "non-adopters" rather than

laggard adopters. What should be clear, however, is that the slowdown of site growth creates pressures regarding who controls the fruits of the community's labor. As one participant in the participatory ratings site (and power-law hub) Yelp.com explained regarding her superuser "elite" status," "It makes you feel special for about two weeks. Then you either realize you're working for someone else without getting paid, you totally lose interest, or you get really into it" (Zittrain, 2008, p. 146).

Sites that fail to effectively manage these governance challenges are at serious risk of "code forking," the term Stephen Weber uses to describe subsets of the open source software community who break off from a large project to start their own similar endeavor (Weber, 2004). Code forking is not inherently a negative attribute—for certain types of communities, particularly those whose collaborations are not particularly granular or modular, there exists a "crowding threshold" above which additional members detract from the community. (See Ciffollili, 2003, for a related discussion of Club Goods theory.). Too much code forking can reduce the network externalities produced by the community, and if one of these forked communities successfully passes the critical mass point in phase 2, then it begins to present a viable alternative to community members who become disaffected over the governance controversies. Likewise, the community must deal with these governance challenges while also embracing new code-based innovations; otherwise it runs the risk of being displaced by a new entrant that suddenly offers community members an augmented set of opportunities.

> H7 ("Code forking the public sphere"): Once the phases of rapid growth have concluded, communities associated with "open source politics" will engage in a series of "code forks."

If H7 is correct, it provides a counterpoint of sorts to Hindman's suggestion that power-law hubs reduce the democratizing effects of the Internet to "mythological" status. For while there can be only one central information hub, audience share is not the only goal of online citizens. If networked publics are able to successfully code fork, and in so doing enable a greater number of small- or medium-sized groups to efficaciously engage in collective action, then the limitations of power-law topology prove a good deal less limiting.

CONCLUSION

The central purpose of this article has been to derive a set of hypotheses about the currently ill-defined field of "open source politics" from the substantial commons-based peer production success represented by Wikipedia.org. In the process of developing these hypotheses, it has been necessary to elaborate several core concepts about Internet-mediated communication that frequently are misunderstood, and also to look at how the standard diffusion-of-innovations process influences the institutional development of online communities.

A key component of the argument lies in the reinterpretation of what power laws in online traffic actually represent. Whereas previous scholars have taken the existence of power laws to indicate a stark limitation on the democratizing impact of the medium, this paper argues that "power-law topology" is of critical importance in solving the collective action problem under conditions of abundance. Whereas traditional research on collective action and public goods problems has focused on the free rider problem, it is self-evident that Wikipedia has no *problem* with free ridership per se. Rather, the problem with online communities of this type lies in coordinating activity in an anti-geographic landscape. Preferential attachment serves as a solution to that challenge, and when combined with the growth inherent in a diffusion process, preferential attachment produces power law distributions. Power laws then serve to solve the anti-geographic challenge presented by the Internet's abundance.

The article concludes with a series of hypotheses specifically because it is my contention that the field of "open source politics" is still in its formative stages. Open source is a product, a process, and a community. Many

have focused on the process and product, while community-formation has gone largely underexplored. This has led to the promulgation of several misperceptions, particularly within the practitioner community and, to a lesser extent, among academic observers. The hypotheses in this article are meant as a challenge of sorts, both to their author and to the research community as a whole. Web research faces the challenge and the opportunity of having massive amounts of data, and much of it noisy. The seven hypotheses are meant to inform our discussion of what open source as a product and process will look like in the political sphere, given the institutional challenges all online communities face as they move through the stages of diffusion. Commons-based peer production processes are not a "field of dreams," capable of being assembled in any situation. They face the limits of existing interest horizons, as determined by early adopter communities rather than the smaller, more focused lead adopter communities. They can only be applied to joint outcomes that prove amenable to online reputation tracking, though the sphere of such outcomes is likely growing apace with the development of the mobile Web. Existing power law hubs are far from immutable due to the constant innovation made possible by Moore's Law, and they thus exhibit the properties of punctuated equilibrium, with long periods of stability and brief shakeups. As an online community concerned with matters of the public sphere matures, and growth slows, this leads to a *healthy* equivalent to the "code forking" discussed by Weber, Kelty, and others. As a result, "open source politics" should not be expected to radically transform the public sphere, but it should render the elite system more porous.

NOTES

1. According to Alexa.com, November 3, 2010.
2. According to Alexa.com, November 3, 2010.
3. According to http://en.wikipedia.org/wiki/Wikipedia:About as of November 3, 2010.
4. Nupedia, unlike Wikipedia, was designed as a for-profit venture of Wales's company, Bomis.com. While entries were to be free, the site was intended to generate revenue through ad sales. Wikipedia was eventually launched as a separate nonprofit after a controversy among volunteer "Wikipedians" over whether the company would one day profit from their free labor.
5. Today, the three rules have been expanded to "five pillars:" (1) Wikipedia is an encyclopedia. (2) Wikipedia has a neutral point of view. (3) Wikipedia is free content. (4) Wikipedia has a code of conduct. (5) Wikipedia does not have firm rules.
6. http://alexa.com/data/details/traffic_details/wikipedia.org. This accords with the findings from the 2007 Pew survey, but tracks traffic on the global, rather than national, level.
7. Olson suggests that two types of groups will face minimal free rider problems. Small groups will be able to identify noncontributors, creating reputation pressures and incentives to recognizably participate (Chong, 1991, develops this case further with regards to social movements). Privileged groups feature a single wealthy participant who will provide as much of the public good as he or she likes regardless. If the wealthy patron has a strong enough taste for the good, all will be satisfied regardless of free riding.
8. http://en.citizendium.org/wiki/CZ:About. Also see http://alexa.com/siteinfo/wikipedia.org+citizendium.org.
9. Obtained through Google Scholar: http://scholar.google.com/scholar?hl=en&lr=&cites=7511022991152445218.
10. See http://radar.oreilly.com/2008/10/network-effects-in-data.html.
11. Research presentation at the Oxford Internet Institute, March 2009. Available online at www.oii.ox.ac.uk.
12. This last point presents a host of measurement issues for social scientists interested in the Internet. User accounts, once created, are rarely destroyed. Competition between social networks, community blogs, or Internet-mediated political associations must be measured in *activity*, rather than list size. Unfortunately, activity measures are almost universally proprietary data, when they are available at all.
13. Image obtained through Google images: http://images.google.com/imgres?imgurl=http://www.cyfm.net/articles/images/S-CurveDetail.jpg&imgrefurl=http://www.cyfm.net/article.php%3Farticle%3DDont_Good_Ideas_Fly.html&h=900&w=900&sz=67&hl=en&start=4&sig2=aXHLBuRCvt8cFz6sdhO5Ag&tbnid=cAFVuhipoW-dfM:&tbnh=146&tbnw=146&ei=trjzR6K_FYyGeufesIgB&prev=/images%3Fq%3Ds-curve%2B%26hl%3Den%26lr%3D%26safe%3Doff%26sa%3DG.
14. Data compiled for "History of Wikipedia" by user "TakuyaMurata." Image available at http://en.wikipedia.org/wiki/History_of_Wikipedia under Creative Commons license.

REFERENCES

Albert, R., Jeong, H., & Barabasi, A. L. (1999). Diameter of the World Wide Web. *Nature*, (401), 130–131.

Barabasi, A. L. (2003). *Linked*. New York, NY: Plume Books.

Baumgartner, F., & Jones, B. (1993). *Agendas and instability in American politics*. Chicago, IL: University of Chicago Press.

Benkler, Y. (2006). *The wealth of networks*. New Haven, CT: Yale University Press.

Bimber, B., Flanagin, A., & Stohl, C. (2005). Reconceptualizing collective action in the contemporary media environment. *Communication Theory*, *15*(4), 365–388.

boyd, d. (2006, March 21). *Friendster lost steam. Is MySpace just a fad?* Blog post. Retrieved from http://www.danah.org/papers/FriendsterMySpace Essay.html

boyd, d. (2009, June). *The not-so-hidden politics of class online*. Presented at the Personal Democracy Forum, New York, NY.

boyd, d., & Ellison, N. (2007). Social network sites: Definition, history, and scholarship. *Journal of Computer Mediated Communication*, *13*(1), article 11.

Bruns, A. (2008). *Blogs, Wikipedia, second life, and beyond: From production to produsage*. New York, NY: Peter Lang.

Chong, D. (1991). *Collective action and the civil rights movement*. Chicago, IL: University of Chicago Press.

Ciffolilli, A. (2003). Phantom authority, self-selective recruitment and retention of members in virtual communities: The case of Wikipedia. *First Monday*, *8*(12). Retrieved from http://firstmonday.org/issues/issue8_12/ciffolilli

Drezner, D., & Farrell, H. (2008). The power and politics of blogs. *Public Choice*, *134*(1–2), 15–30.

Fine, A., Sifry, M., Raseij, A., & Levy, J. (Eds.). (2008). *Rebooting America: Ideas for redesigning American democracy for the Internet age*. New York, NY: Personal Democracy Press.

Giles, J. (2005). Internet encyclopedias go head to head. *Nature*, *438*, 900–901.

Hara, N. (2008). Internet use for political mobilization: Voices of participants. *First Monday*, *13*(7).

Hindman, M. (2007). Open-source politics reconsidered: Emerging patterns in online political participation. In V. Mayer-Schonberger and D. Lazer (Eds.), *Governance and information technology: From electronic government to information government* (pp. 183–207). Cambridge, MA: MIT Press.

Hindman, M. (2008). *The myth of digital democracy*. Princeton, NJ: Princeton University Press.

Hindman, M., Tsioutsiouliklis, K., & Johnson, J. (2003, April). "*Googlearchy": How a few heavily-linked sites dominate politics on the Web*. Paper presented at the Annual Meeting of the Midwest Political Science Association, Chicago, IL.

Huberman, B., & Wilkinson, D. (2007). Assessing the value of cooperation in Wikipedia. *First Monday, 12*(4). Retrieved from http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/ 1763/1643

Jenkins, H. (2006). *Convergence culture: Where old and new media collide*. New York, NY: NYU Press.

Karpf, D. (2008, March). Measuring influence in the political blogosphere: Who's winning, and how can we tell? *Politics and Technology Review*, 33–41. Retrieved from http://www.asiaing.com/politics-and-technology-review-march-2008.html

Karpf, D. (2009a, September). *The MoveOn effect: Disruptive innovation within the interest group ecology of American politics*. Presented at American Political Science Association Annual Meeting, Toronto, Ontario, Canada.

Karpf, D. (2009b, October). *Don't think of an online elephant: Explaining the dearth of conservative online political infrastructure in America*. Presented at Society for the Social Study of the Sciences Annual Meeting, Washington, DC.

Karpf, D. (2010). Macaca moments reconsidered: Electoral panopticon or Netroots mobilization? *Journal of Information Technology & Politics*, *7*(2–3), 143–162.

Karpf, D. (2011). Why bowl alone when you can flash-mob the bowling alley? *IEEE Intelligent Systems,* 26(1), 40–47.

Kelty, C. (2008). *Two bits: The cultural significance of free software*. Durham, NC: Duke University Press.

Kingdon, J. (1984). *Agendas, alternatives, and public policy*. New York, NY: Longman Publishers.

Kollock, P. (1999). The economies of online cooperation: Gifts and public goods in cyberspace. In M. A. Smith and P. Kollock (Eds.), *Communities in cyberspace* (pp. 220–239). New York, NY: Routledge.

Lessig, L. (1999). *Code and other laws of cyberspace*. New York, NY: Basic Books.

Lessig, L. (2003, August 19). Interview with Joe Trippi. *Lessig.org*, blog entry. Retrieved from http://www.lessig.org/blog/archives/001428.shtml

Lev-On, A., & Hardin, R. (2008). Internet-based collaborations and their political significance. *Journal of Information Technology & Politics*, *4*(2), 5–27.

Lih, A. (2009). *The Wikipedia revolution*. New York, NY: Hyperion Books.

Lupia, A., & Sin, G. (2003). Which public goods are endangered? How evolving communication technologies affect the logic of collective action. *Public Choice*, *117*(3–4), 315–331.

Moore, G. (1965). Cramming more components onto integrated circuits. *Electronics Magazine*, 4, 114–117.

Olson, M. (1965). *The logic of collective action*. Cambridge, MA: Harvard University Press.

Rainie, L., & Tancer, B. (2007). *Wikipedia users*. Research Report. Washington, DC: Pew Internet & American Life Project. Retrieved from http://www.pewinternet.org/PPF/r/212/report_display.asp

Ratcliffe, M., & Lebkowsky, J., Eds. (2005). *Extreme democracy*. Stanford, CA: Creative Commons.

Raymond, E. S. (2001). *The cathedral & the bazaar: Musings on Linux and open source by an accidental revolutionary*. Sebastopol, CA: O'Reilly.

Reagle, J. (2010). *Good faith collaboration: The culture of Wikipedia*. Cambridge, MA: MIT Press.

Reed, D. (1999). That sneaky exponential—beyond Metcalfe's Law to the power of community building. *Context Magazine* (online supplement). Retrieved from http://www.reed.com/dpr/locus/gfn/reedslaw.html

Resnick, P., Zeckhauser, R., Swanson, J., & Lockwood, K. (2006). The value of reputation on eBay: A controlled experiment. *Experimental Economics*, 9(2), 79–101.

Rogers, E. (2003). *Diffusion of innovation* (5th ed.). New York, NY: Free Press.

Rushkoff, D. (2003). *Open source democracy: How online communication is changing offline politics*. New York, NY: Demos.

Schudson, M. (1999). *The good citizen: A history of American civic life*. New York, NY: Free Press.

Seelye, K. (2005, December 4). Snared in the Web of a Wikipedia Liar. *The New York Times*. Retrieved from http://www.nytimes.com/2005/12/04/weekinreview/04seelye.html

Shirky, C. (1999). *The interest horizons and the limits of software love*. Self-published. Retrieved from http://www.shirky.com/writings/interest.html

Shirky, C. (2005). Power laws, Weblogs, and inequality. In M. Ratcliffe & J. Lebkowsky (Eds.), *Extreme democracy* (pp. 49–57). Stanford, CA: Creative Commons.

Shirky, C. (2008). *Here comes everybody*. New York, NY: Penguin Press.

Sifry, M. (2004, November 4). The rise of open-source politics. *The Nation*. Retrieved from http://www.thenation.com/doc.mhtml?i=20041122&s=sifry

Stalder, F. (2006). *Manuel Castells and the theory of network society*. New York, NY: Polity Press.

Trippi, J. (2004). *The revolution will not be televised*. New York, NY: Regan Books.

Valente, T. (1995). *Network models of diffusion of innovation*. New York, NY: Hampton Press.

Viegas, F., Wattenberg, M., Kriss, J., & Van Ham, F. (2007). Talk before you type: Coordination in Wikipedia. *System Sciences, 2007*. Proceedings of the 40th Annual Hawaii International Conference on System Sciences. Retrieved from http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=4076527

Vise, D., & Malseed, M. (2005). *The Google story: Inside the hottest business and technology success of our time*. New York, NY: Delacorte Press.

Von Hippel, E. (2005). *Democratizing innovation*. Cambridge, MA: MIT Press.

Wasserman, S., & Faust, K. (1994). *Social network analysis: Methods and applications*. New York, NY: Cambridge University Press.

Weber, S. (2004). *The success of open source*. Cambridge, MA: Harvard University Press.

Weinberger, D. (2007). *Everything is miscellaneous*. New York, NY: Holt Paperbacks.

Zittrain, J. (2008). *The future of the Internet (and how to stop it)*. New Haven, CT: Yale University Press.