

Chapter 7 Political Freedom Part 2: Emergence of the Networked Public Sphere

The fundamental elements of the difference between the networked information economy and the mass media are network architecture and the cost of becoming a speaker. The first element is the shift from a hub-and-spoke architecture with unidirectional links to the end points in the mass media, to distributed architecture with multidirectional connections among all nodes in the networked information environment. The second is the practical elimination of communications costs as a barrier to speaking across associational boundaries. Together, these characteristics have fundamentally altered the capacity of individuals, acting alone or with others, to be active participants in the public sphere as opposed to its passive readers, listeners, or viewers. For authoritarian countries, this means that it is harder and more costly, though not perhaps entirely impossible, to both be networked and maintain control over their public spheres. China seems to be doing too good a job of this in the middle of the first decade of this century for us to say much more than that it is harder to maintain control, and therefore that at least in some authoritarian regimes, control will be looser. In

liberal democracies, ubiquitous individual ability to produce information creates the potential for near-universal intake. It therefore portends significant, though not inevitable, changes in the structure of the public sphere from the commercial mass-media environment. These changes raise challenges for filtering. They underlie some of the critiques of the claims about the democratizing effect of the Internet that I explore later in this chapter. Fundamentally, however, they are the roots of possible change. Beginning with the cost of sending an e-mail to some number of friends or to a mailing list of people interested in a particular subject, to the cost of setting up a Web site or a blog, and through to the possibility of maintaining interactive conversations with large numbers of people through sites like Slashdot, the cost of being a speaker in a regional, national, or even international political conversation is several orders of magnitude lower than the cost of speaking in the mass-mediated environment. This, in turn, leads to several orders of magnitude more speakers and participants in conversation and, ultimately, in the public sphere.

The change is as much qualitative as it is quantitative. The qualitative change is represented in the experience of being a potential speaker, as opposed to simply a listener and voter. It relates to the self-perception of individuals in society and the culture of participation they can adopt. The easy possibility of communicating effectively into the public sphere allows individuals to reorient themselves from passive readers and listeners to potential speakers and participants in a conversation. The way we listen to what we hear changes because of this; as does, perhaps most fundamentally, the way we observe and process daily events in our lives. We no longer need to take these as merely private observations, but as potential subjects for public communication. This change affects the relative power of the media. It affects the structure of intake of observations and views. It affects the presentation of issues and observations for discourse. It affects the way issues are filtered, for whom and by whom. Finally, it affects the ways in which positions are crystallized and synthesized, sometimes still by being amplified to the point that the mass media take them as inputs and convert them into political positions, but occasionally by direct organization of opinion and action to the point of reaching a salience that drives the political process directly.

The basic case for the democratizing effect of the Internet, as seen from the perspective of the mid-1990s, was articulated in an opinion of the U.S. Supreme Court in *Reno v. ACLU*:

The Web is thus comparable, from the readers' viewpoint, to both a vast library including millions of readily available and indexed publications and a sprawling mall offering goods and services. From the publishers' point of view, it constitutes a vast platform from which to address and hear from a world-wide audience of millions of readers, viewers, researchers, and buyers. Any person or organization with a computer connected to the Internet can "publish" information. Publishers include government agencies, educational institutions, commercial entities, advocacy groups, and individuals. . . .

Through the use of chat rooms, any person with a phone line can become a town crier with a voice that resonates farther than it could from any soapbox. Through the use of Web pages, mail exploders, and newsgroups, the same individual can become a pamphleteer. As the District Court found, "the content on the Internet is as diverse as human thought."¹

The observations of what is different and unique about this new medium relative to those that dominated the twentieth century are already present in the quotes from the Court. There are two distinct types of effects. The first, as the Court notes from "the readers' perspective," is the abundance and diversity of human expression available to anyone, anywhere, in a way that was not feasible in the mass-mediated environment. The second, and more fundamental, is that anyone can be a publisher, including individuals, educational institutions, and nongovernmental organizations (NGOs), alongside the traditional speakers of the mass-media environment—government and commercial entities.

Since the end of the 1990s there has been significant criticism of this early conception of the democratizing effects of the Internet. One line of critique includes variants of the Babel objection: the concern that information overload will lead to fragmentation of discourse, polarization, and the loss of political community. A different and descriptively contradictory line of critique suggests that the Internet is, in fact, exhibiting concentration: Both infrastructure and, more fundamentally, patterns of attention are much less distributed than we thought. As a consequence, the Internet diverges from the mass media much less than we thought in the 1990s and significantly less than we might hope.

I begin the chapter by offering a menu of the core technologies and usage patterns that can be said, as of the middle of the first decade of the twenty-first century, to represent the core Internet-based technologies of democratic discourse. I then use two case studies to describe the social and economic practices through which these tools are implemented to construct the public

sphere, and how these practices differ quite radically from the mass-media model. On the background of these stories, we are then able to consider the critiques that have been leveled against the claim that the Internet democratizes. Close examination of the application of networked information economy to the production of the public sphere suggests that the emerging networked public sphere offers significant improvements over one dominated by commercial mass media. Throughout the discussion, it is important to keep in mind that the relevant comparison is always between the public sphere that we in fact had throughout the twentieth century, the one dominated by mass media, that is the baseline for comparison, not the utopian image of the "everyone a pamphleteer" that animated the hopes of the 1990s for Internet democracy. Departures from the naïve utopia are not signs that the Internet does not democratize, after all. They are merely signs that the medium and its analysis are maturing.

BASIC TOOLS OF NETWORKED COMMUNICATION

Analyzing the effect of the networked information environment on public discourse by cataloging the currently popular tools for communication is, to some extent, self-defeating. These will undoubtedly be supplanted by new ones. Analyzing this effect without having a sense of what these tools are or how they are being used is, on the other hand, impossible. This leaves us with the need to catalog what is, while trying to abstract from what is being used to what relationships of information and communication are emerging, and from these to transpose to a theory of the networked information economy as a new platform for the public sphere.

E-mail is the most popular application on the Net. It is cheap and trivially easy to use. Basic e-mail, as currently used, is not ideal for public communications. While it provides a cheap and efficient means of communicating with large numbers of individuals who are not part of one's basic set of social associations, the presence of large amounts of commercial spam and the amount of mail flowing in and out of mailboxes make indiscriminate e-mail distributions a relatively poor mechanism for being heard. E-mails to smaller groups, preselected by the sender for having some interest in a subject or relationship to the sender, do, however, provide a rudimentary mechanism for communicating observations, ideas, and opinions to a significant circle, on an ad hoc basis. Mailing lists are more stable and self-selecting, and

therefore more significant as a basic tool for the networked public sphere. Some mailing lists are moderated or edited, and run by one or a small number of editors. Others are not edited in any significant way. What separates mailing lists from most Web-based uses is the fact that they push the information on them into the mailbox of subscribers. Because of their attention limits, individuals restrict their subscriptions, so posting on a mailing list tends to be done by and for people who have self-selected as having a heightened degree of common interest, substantive or contextual. It therefore enhances the degree to which one is heard by those already interested in a topic. It is not a communications model of one-to-many, or few-to-many as broadcast is to an open, undefined class of audience members. Instead, it allows one, or a few, or even a limited large group to communicate to a large but limited group, where the limit is self-selection as being interested or even immersed in a subject.

The World Wide Web is the other major platform for tools that individuals use to communicate in the networked public sphere. It enables a wide range of applications, from basic static Web pages, to, more recently, blogs and various social-software-mediated platforms for large-scale conversations of the type described in chapter 3—like Slashdot. Static Web pages are the individual's basic "broadcast" medium. They allow any individual or organization to present basic texts, sounds, and images pertaining to their position. They allow small NGOs to have a worldwide presence and visibility. They allow individuals to offer thoughts and commentaries. They allow the creation of a vast, searchable database of information, observations, and opinions, available at low cost for anyone, both to read and write into. This does not yet mean that all these statements are heard by the relevant others to whom they are addressed. Substantial analysis is devoted to that problem, but first let us complete the catalog of tools and information flow structures.

One Web-based tool and an emerging cultural practice around it that extends the basic characteristics of Web sites as media for the political public sphere are Web logs, or blogs. Blogs are a tool and an approach to using the Web that extends the use of Web pages in two significant ways. Technically, blogs are part of a broader category of innovations that make the web "writable." That is, they make Web pages easily capable of modification through a simple interface. They can be modified from anywhere with a networked computer, and the results of writing onto the Web page are immediately available to anyone who accesses the blog to read. This technical change resulted in two divergences from the cultural practice of Web sites

in the 1990s. First, they allowed the evolution of a journal-style Web page, where individual short posts are added to the Web site in short or large intervals. As practice has developed over the past few years, these posts are usually archived chronologically. For many users, this means that blogs have become a form of personal journal, updated daily or so, for their own use and perhaps for the use of a very small group of friends. What is significant about this characteristic from the perspective of the construction of the public sphere is that blogs enable individuals to write to their Web pages in journalism time—that is, hourly, daily, weekly—whereas Web page culture that preceded it tended to be slower moving: less an equivalent of reportage than of the essay. Today, one certainly finds individuals using blog software to maintain what are essentially static Web pages, to which they add essays or content occasionally, and Web sites that do not use blogging technology but are updated daily. The public sphere function is based on the content and cadence—that is, the use practice—not the technical platform.

The second critical innovation of the writable Web in general and of blogs in particular was the fact that in addition to the owner, readers/users could write to the blog. Blogging software allows the person who runs a blog to permit some, all, or none of the readers to post comments to the blog, with or without retaining power to edit or moderate the posts that go on, and those that do not. The result is therefore not only that many more people write finished statements and disseminate them widely, but also that the end product is a weighted conversation, rather than a finished good. It is a conversation because of the common practice of allowing and posting comments, as well as comments to these comments. Blog writers—bloggers—often post their own responses in the comment section or address comments in the primary section. Blog-based conversation is weighted, because the culture and technical affordances of blogging give the owner of the blog greater weight in deciding who gets to post or comment and who gets to decide these questions. Different blogs use these capabilities differently; some opt for broader intake and discussion on the board, others for a more tightly edited blog. In all these cases, however, the communications model or information-flow structure that blogs facilitate is a weighted conversation that takes the form of one or a group of primary contributors/authors, together with some larger number, often many, secondary contributors, communicating to an unlimited number of many readers.

The writable Web also encompasses another set of practices that are distinct, but that are often pooled in the literature together with blogs. These

are the various larger-scale, collaborative-content production systems available on the Web, of the type described in chapter 3. Two basic characteristics make sites like Slashdot or Wikipedia different from blogs. First, they are intended for, and used by, very large groups, rather than intended to facilitate a conversation weighted toward one or a small number of primary speakers. Unlike blogs, they are not media for individual or small group expression with a conversation feature. They are intrinsically group communication media. They therefore incorporate social software solutions to avoid deterioration into chaos—peer review, structured posting privileges, reputation systems, and so on. Second, in the case of Wikis, the conversation platform is anchored by a common text. From the perspective of facilitating the synthesis of positions and opinions, the presence of collaborative authorship of texts offers an additional degree of viscosity to the conversation, so that views “stick” to each other, must jostle for space, and accommodate each other. In the process, the output is more easily recognizable as a collective output and a salient opinion or observation than where the form of the conversation is more free-flowing exchange of competing views.

Common to all these Web-based tools—both static and dynamic, individual and cooperative—are linking, quotation, and presentation. It is at the very core of the hypertext markup language (HTML) to make referencing easy. And it is at the very core of a radically distributed network to allow materials to be archived by whoever wants to archive them, and then to be accessible to whoever has the reference. Around these easy capabilities, the cultural practice has emerged to reference through links for easy transition from your own page or post to the one you are referring to—whether as inspiration or in disagreement. This culture is fundamentally different from the mass-media culture, where sending a five-hundred-page report to millions of users is hard and expensive. In the mass media, therefore, instead of allowing readers to read the report alongside its review, all that is offered is the professional review in the context of a culture that trusts the reviewer. On the Web, linking to original materials and references is considered a core characteristic of communication. The culture is oriented toward “see for yourself.” Confidence in an observation comes from a combination of the reputation of the speaker as it has emerged over time, reading underlying sources you believe you have some competence to evaluate for yourself, and knowing that for any given referenced claim or source, there is some group of people out there, unaffiliated with the reviewer or speaker, who will have access to the source and the means for making their disagreement with the

speaker's views known. Linking and "see for yourself" represent a radically different and more participatory model of accreditation than typified the mass media.

Another dimension that is less well developed in the United States than it is in Europe and East Asia is mobility, or the spatial and temporal ubiquity of basic tools for observing and commenting on the world we inhabit. Dan Gillmor is clearly right to include these basic characteristics in his book *We the Media*, adding to the core tools of what he describes as a transformation in journalism, short message service (SMS), and mobile connected cameras to mailing lists, Web logs, Wikis, and other tools. The United States has remained mostly a PC-based networked system, whereas in Europe and Asia, there has been more substantial growth in handheld devices, primarily mobile phones. In these domains, SMS—the "e-mail" of mobile phones—and camera phones have become critical sources of information, in real time. In some poor countries, where cell phone minutes remain very (even prohibitively) expensive for many users and where landlines may not exist, text messaging is becoming a central and ubiquitous communication tool. What these suggest to us is a transition, as the capabilities of both systems converge, to widespread availability of the ability to register and communicate observations in text, audio, and video, wherever we are and whenever we wish. Drazen Pantic tells of how listeners of Internet-based Radio B-92 in Belgrade reported events in their neighborhoods after the broadcast station had been shut down by the Milosevic regime. Howard Rheingold describes in *Smart Mobs* how citizens of the Philippines used SMS to organize real-time movements and action to overthrow their government. In a complex modern society, where things that matter can happen anywhere and at any time, the capacities of people armed with the means of recording, rendering, and communicating their observations change their relationship to the events that surround them. Whatever one sees and hears can be treated as input into public debate in ways that were impossible when capturing, rendering, and communicating were facilities reserved to a handful of organizations and a few thousands of their employees.

NETWORKED INFORMATION ECONOMY MEETS THE PUBLIC SPHERE

The networked public sphere is not made of tools, but of social production practices that these tools enable. The primary effect of the Internet on the

public sphere in liberal societies relies on the information and cultural production activity of emerging nonmarket actors: individuals working alone and cooperatively with others, more formal associations like NGOs, and their feedback effect on the mainstream media itself. These enable the networked public sphere to moderate the two major concerns with commercial mass media as a platform for the public sphere: (1) the excessive power it gives its owners, and (2) its tendency, when owners do not dedicate their media to exert power, to foster an inert polity. More fundamentally, the social practices of information and discourse allow a very large number of actors to see themselves as potential contributors to public discourse and as potential actors in political arenas, rather than mostly passive recipients of mediated information who occasionally can vote their preferences. In this section, I offer two detailed stories that highlight different aspects of the effects of the networked information economy on the construction of the public sphere. The first story focuses on how the networked public sphere allows individuals to monitor and disrupt the use of mass-media power, as well as organize for political action. The second emphasizes in particular how the networked public sphere allows individuals and groups of intense political engagement to report, comment, and generally play the role traditionally assigned to the press in observing, analyzing, and creating political salience for matters of public interest. The case studies provide a context both for seeing how the networked public sphere responds to the core failings of the commercial, mass-media-dominated public sphere and for considering the critiques of the Internet as a platform for a liberal public sphere.

Our first story concerns Sinclair Broadcasting and the 2004 U.S. presidential election. It highlights the opportunities that mass-media owners have to exert power over the public sphere, the variability within the media itself in how this power is used, and, most significant for our purposes here, the potential corrective effect of the networked information environment. At its core, it suggests that the existence of radically decentralized outlets for individuals and groups can provide a check on the excessive power that media owners were able to exercise in the industrial information economy.

Sinclair, which owns major television stations in a number of what were considered the most competitive and important states in the 2004 election—including Ohio, Florida, Wisconsin, and Iowa—informed its staff and stations that it planned to preempt the normal schedule of its sixty-two stations to air a documentary called *Stolen Honor: The Wounds That Never Heal*, as a news program, a week and a half before the elections.² The documentary

was reported to be a strident attack on Democratic candidate John Kerry's Vietnam War service. One reporter in Sinclair's Washington bureau, who objected to the program and described it as "blatant political propaganda," was promptly fired.³ The fact that Sinclair owns stations reaching one quarter of U.S. households, that it used its ownership to preempt local broadcast schedules, and that it fired a reporter who objected to its decision, make this a classic "Berlusconi effect" story, coupled with a poster-child case against media concentration and the ownership of more than a small number of outlets by any single owner. The story of Sinclair's plans broke on Saturday, October 9, 2004, in the *Los Angeles Times*. Over the weekend, "official" responses were beginning to emerge in the Democratic Party. The Kerry campaign raised questions about whether the program violated election laws as an undeclared "in-kind" contribution to the Bush campaign. By Tuesday, October 12, the Democratic National Committee announced that it was filing a complaint with the Federal Elections Commission (FEC), while seventeen Democratic senators wrote a letter to the chairman of the Federal Communications Commission (FCC), demanding that the commission investigate whether Sinclair was abusing the public trust in the airwaves. Neither the FEC nor the FCC, however, acted or intervened throughout the episode.

Alongside these standard avenues of response in the traditional public sphere of commercial mass media, their regulators, and established parties, a very different kind of response was brewing on the Net, in the blogosphere. On the morning of October 9, 2004, the *Los Angeles Times* story was blogged on a number of political blogs—Josh Marshall on talkingpointsmemo.com, Chris Bower on MyDD.com, and Markos Moulitsas on dailyKos.com. By midday that Saturday, October 9, two efforts aimed at organizing opposition to Sinclair were posted in the dailyKos and MyDD. A "boycott-Sinclair" site was set up by one individual, and was pointed to by these blogs. Chris Bowers on MyDD provided a complete list of Sinclair stations and urged people to call the stations and threaten to picket and boycott. By Sunday, October 10, the dailyKos posted a list of national advertisers with Sinclair, urging readers to call them. On Monday, October 11, MyDD linked to that list, while another blog, theleftcoaster.com, posted a variety of action agenda items, from picketing affiliates of Sinclair to suggesting that readers oppose Sinclair license renewals, providing a link to the FCC site explaining the basic renewal process and listing public-interest organizations to work with. That same day, another individual, Nick Davis, started a Web site,

BoycottSBG.com, on which he posted the basic idea that a concerted boycott of local advertisers was the way to go, while another site, stopsinclair.org, began pushing for a petition. In the meantime, TalkingPoints published a letter from Reed Hundt, former chairman of the FCC, to Sinclair, and continued finding tidbits about the film and its maker. Later on Monday, TalkingPoints posted a letter from a reader who suggested that stockholders of Sinclair could bring a derivative action. By 5:00 A.M. on the dawn of Tuesday, October 12, however, TalkingPoints began pointing toward Davis's database on BoycottSBG.com. By 10:00 that morning, Marshall posted on TalkingPoints a letter from an anonymous reader, which began by saying: "I've worked in the media business for 30 years and I guarantee you that sales is what these local TV stations are all about. They don't care about license renewal or overwhelming public outrage. They care about sales only, so only local advertisers can affect their decisions." This reader then outlined a plan for how to watch and list all local advertisers, and then write to the sales managers—not general managers—of the local stations and tell them which advertisers you are going to call, and then call those. By 1:00 P.M. Marshall posted a story of his own experience with this strategy. He used Davis's database to identify an Ohio affiliate's local advertisers. He tried to call the sales manager of the station, but could not get through. He then called the advertisers. The post is a "how to" instruction manual, including admonitions to remember that the advertisers know nothing of this, the story must be explained, and accusatory tones avoided, and so on. Marshall then began to post letters from readers who explained with whom they had talked—a particular sales manager, for example—and who were then referred to national headquarters. He continued to emphasize that advertisers were the right addressees. By 5:00 P.M. that same Tuesday, Marshall was reporting more readers writing in about experiences, and continued to steer his readers to sites that helped them to identify their local affiliate's sales manager and their advertisers.⁴

By the morning of Wednesday, October 13, the boycott database already included eight hundred advertisers, and was providing sample letters for users to send to advertisers. Later that day, BoycottSBG reported that some participants in the boycott had received reply e-mails telling them that their unsolicited e-mail constituted illegal spam. Davis explained that the CAN-SPAM Act, the relevant federal statute, applied only to commercial spam, and pointed users to a law firm site that provided an overview of CAN-SPAM. By October 14, the boycott effort was clearly bearing fruit. Davis

reported that Sinclair affiliates were threatening advertisers who cancelled advertisements with legal action, and called for volunteer lawyers to help respond. Within a brief period, he collected more than a dozen volunteers to help the advertisers. Later that day, another blogger at grassrootsnation.com had set up a utility that allowed users to send an e-mail to all advertisers in the BoycottSBG database. By the morning of Friday, October 15, Davis was reporting more than fifty advertisers pulling ads, and three or four mainstream media reports had picked up the boycott story and reported on it. That day, an analyst at Lehman Brothers issued a research report that downgraded the expected twelve-month outlook for the price of Sinclair stock, citing concerns about loss of advertiser revenue and risk of tighter regulation. Mainstream news reports over the weekend and the following week systematically placed that report in context of local advertisers pulling their ads from Sinclair. On Monday, October 18, the company's stock price dropped by 8 percent (while the S&P 500 rose by about half a percent). The following morning, the stock dropped a further 6 percent, before beginning to climb back, as Sinclair announced that it would not show *Stolen Honor*, but would provide a balanced program with only portions of the documentary and one that would include arguments on the other side. On that day, the company's stock price had reached its lowest point in three years. The day after the announced change in programming decision, the share price bounced back to where it had been on October 15. There were obviously multiple reasons for the stock price losses, and Sinclair stock had been losing ground for many months prior to these events. Nonetheless, as figure 7.1 demonstrates, the market responded quite sluggishly to the announcements of regulatory and political action by the Democratic establishment earlier in the week of October 12, by comparison to the precipitous decline and dramatic bounce-back surrounding the market projections that referred to advertising loss. While this does not prove that the Web-organized, blog-driven and -facilitated boycott was the determining factor, as compared to fears of formal regulatory action, the timing strongly suggests that the efficacy of the boycott played a very significant role.

The first lesson of the Sinclair *Stolen Honor* story is about commercial mass media themselves. The potential for the exercise of inordinate power by media owners is not an imaginary concern. Here was a publicly traded firm whose managers supported a political party and who planned to use their corporate control over stations reaching one quarter of U.S. households, many in swing states, to put a distinctly political message in front of this

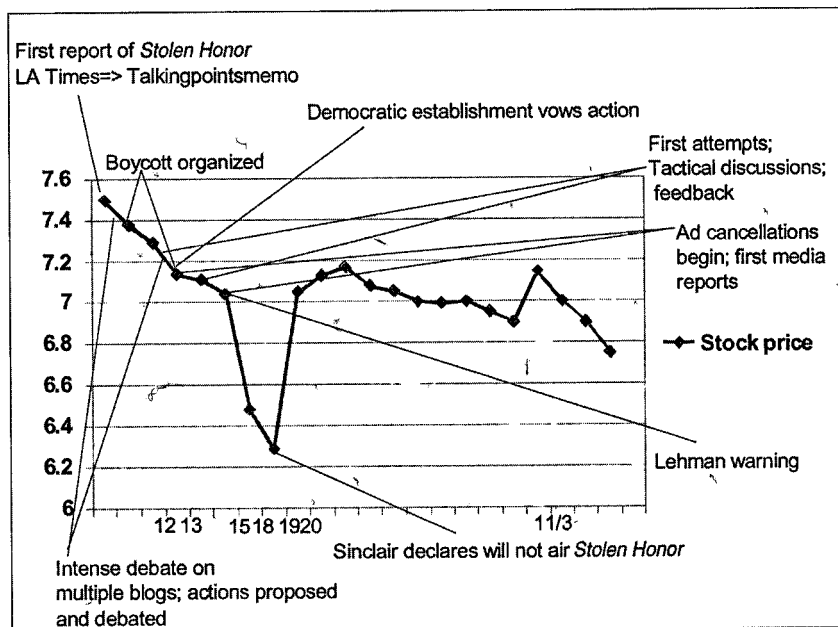


Figure 7.1: Sinclair Stock, October 8–November 5, 2004

large audience. We also learn, however, that in the absence of monopoly, such decisions do not determine what everyone sees or hears, and that other mass-media outlets will criticize each other under these conditions. This criticism alone, however, cannot stop a determined media owner from trying to exert its influence in the public sphere, and if placed as Sinclair was, in locations with significant political weight, such intervention could have substantial influence. Second, we learn that the new, network-based media can exert a significant counterforce. They offer a completely new and much more widely open intake basin for insight and commentary. The speed with which individuals were able to set up sites to stake out a position, to collect and make available information relevant to a specific matter of public concern, and to provide a platform for others to exchange views about the appropriate political strategy and tactics was completely different from anything that the economics and organizational structure of mass media make feasible. The third lesson is about the internal dynamics of the networked public sphere. Filtering and synthesis occurred through discussion, trial, and error. Multiple proposals for action surfaced, and the practice of linking allowed most anyone interested who connected to one of the nodes in the network to follow

quotations and references to get a sense of the broad range of proposals. Different people could coalesce on different modes of action—150,000 signed the petition on stopsinclair.org, while others began to work on the boycott. Setting up the mechanism was trivial, both technically and as a matter of cost—something a single committed individual could choose to do. Pointing and adoption provided the filtering, and feedback about the efficacy, again distributed through a system of cross-references, allowed for testing and accreditation of this course of action. High-visibility sites, like Talkingpointsmemo or the dailyKos, offered transmissions hubs that disseminated information about the various efforts and provided a platform for interest-group-wide tactical discussions. It remains ambiguous to what extent these dispersed loci of public debate still needed mass-media exposure to achieve broad political salience. BoycottSBG.com received more than three hundred thousand unique visitors during its first week of operations, and more than one million page views. It successfully coordinated a campaign that resulted in real effects on advertisers in a large number of geographically dispersed media markets. In this case, at least, mainstream media reports on these efforts were few, and the most immediate “transmission mechanism” of their effect was the analyst’s report from Lehman, not the media. It is harder to judge the extent to which those few mainstream media reports that did appear featured in the decision of the analyst to credit the success of the boycott efforts. The fact that mainstream media outlets may have played a role in increasing the salience of the boycott does not, however, take away from the basic role played by these new mechanisms of bringing information and experience to bear on a broad public conversation combined with a mechanism to organize political action across many different locations and social contexts.

Our second story focuses not on the new reactive capacity of the networked public sphere, but on its generative capacity. In this capacity, it begins to outline the qualitative change in the role of individuals as potential investigators and commentators, as active participants in defining the agenda and debating action in the public sphere. This story is about Diebold Election Systems (one of the leading manufacturers of electronic voting machines and a subsidiary of one of the foremost ATM manufacturers in the world, with more than \$2 billion a year in revenue), and the way that public criticism of its voting machines developed. It provides a series of observations about how the networked information economy operates, and how it allows large numbers of people to participate in a peer-production enterprise of

news gathering, analysis, and distribution, applied to a quite unsettling set of claims. While the context of the story is a debate over electronic voting, that is not what makes it pertinent to democracy. The debate could have centered on any corporate and government practice that had highly unsettling implications, was difficult to investigate and parse, and was largely ignored by mainstream media. The point is that the networked public sphere did engage, and did successfully turn something that was not a matter of serious public discussion to a public discussion that led to public action.

Electronic voting machines were first used to a substantial degree in the United States in the November 2002 elections. Prior to, and immediately following that election, there was sparse mass-media coverage of electronic voting machines. The emphasis was mostly on the newness, occasional slips, and the availability of technical support staff to help at polls. An *Atlanta Journal-Constitution* story, entitled "Georgia Puts Trust in Electronic Voting, Critics Fret about Absence of Paper Trails,"⁵ is not atypical of coverage at the time, which generally reported criticism by computer engineers, but conveyed an overall soothing message about the efficacy of the machines and about efforts by officials and companies to make sure that all would be well. The *New York Times* report of the Georgia effort did not even mention the critics.⁶ The *Washington Post* reported on the fears of failure with the newness of the machines, but emphasized the extensive efforts that the manufacturer, Diebold, was making to train election officials and to have hundreds of technicians available to respond to failure.⁷ After the election, the *Atlanta Journal-Constitution* reported that the touch-screen machines were a hit, burying in the text any references to machines that highlighted the wrong candidates or the long lines at the booths, while the *Washington Post* highlighted long lines in one Maryland county, but smooth operation elsewhere. Later, the *Post* reported a University of Maryland study that surveyed users and stated that quite a few needed help from election officials, compromising voter privacy.⁸ Given the centrality of voting mechanisms for democracy, the deep concerns that voting irregularities determined the 2000 presidential elections, and the sense that voting machines would be a solution to the "hanging chads" problem (the imperfectly punctured paper ballots that came to symbolize the Florida fiasco during that election), mass-media reports were remarkably devoid of any serious inquiry into how secure and accurate voting machines were, and included a high quotient of soothing comments from election officials who bought the machines and executives of the manufacturers who sold them. No mass-media outlet sought to go

behind the claims of the manufacturers about their machines, to inquire into their security or the integrity of their tallying and transmission mechanisms against vote tampering. No doubt doing so would have been difficult. These systems were protected as trade secrets. State governments charged with certifying the systems were bound to treat what access they had to the inner workings as confidential. Analyzing these systems requires high degrees of expertise in computer security. Getting around these barriers is difficult. However, it turned out to be feasible for a collection of volunteers in various settings and contexts on the Net.

In late January 2003, Bev Harris, an activist focused on electronic voting machines, was doing research on Diebold, which has provided more than 75,000 voting machines in the United States and produced many of the machines used in Brazil's purely electronic voting system. Harris had set up a whistle-blower site as part of a Web site she ran at the time, *blackboxvoting.com*. Apparently working from a tip, Harris found out about an openly available site where Diebold stored more than forty thousand files about how its system works. These included specifications for, and the actual code of, Diebold's machines and vote-tallying system. In early February 2003, Harris published two initial journalistic accounts on an online journal in New Zealand, *Scoop.com*—whose business model includes providing an unedited platform for commentators who wish to use it as a platform to publish their materials. She also set up a space on her Web site for technically literate users to comment on the files she had retrieved. In early July of that year, she published an analysis of the results of the discussions on her site, which pointed out how access to the Diebold open site could have been used to affect the 2002 election results in Georgia (where there had been a tightly contested Senate race). In an editorial attached to the publication, entitled "Bigger than Watergate," the editors of *Scoop* claimed that what Harris had found was nothing short of a mechanism for capturing the U.S. elections process. They then inserted a number of lines that go to the very heart of how the networked information economy can use peer production to play the role of watchdog:

We can now reveal for the first time the location of a complete online copy of the original data set. As we anticipate attempts to prevent the distribution of this information we encourage supporters of democracy to make copies of these files and to make them available on websites and file sharing networks: <http://users.actrix.co.nz/dolly/>. As many of the files are zip password protected you may need some assistance in opening them, we have found that the utility available at

the following URL works well: <http://www.lostpassword.com>. Finally some of the zip files are partially damaged, but these too can be read by using the utility at: <http://www.zip-repair.com/>. At this stage in this inquiry we do not believe that we have come even remotely close to investigating all aspects of this data; i.e., there is no reason to believe that the security flaws discovered so far are the only ones. Therefore we expect many more discoveries to be made. We want the assistance of the online computing community in this enterprise and we encourage you to file your findings at the forum HERE [providing link to forum].

A number of characteristics of this call to arms would have been simply infeasible in the mass-media environment. They represent a genuinely different mind-set about how news and analysis are produced and how censorship and power are circumvented. First, the ubiquity of storage and communications capacity means that public discourse can rely on "see for yourself" rather than on "trust me." The first move, then, is to make the raw materials available for all to see. Second, the editors anticipated that the company would try to suppress the information. Their response was not to use a counterweight of the economic and public muscle of a big media corporation to protect use of the materials. Instead, it was widespread distribution of information—about where the files could be found, and about where tools to crack the passwords and repair bad files could be found—matched with a call for action: get these files, copy them, and store them in many places so they cannot be squelched. Third, the editors did not rely on large sums of money flowing from being a big media organization to hire experts and interns to scour the files. Instead, they posed a challenge to whoever was interested—there are more scoops to be found, this is important for democracy, good hunting!! Finally, they offered a platform for integration of the insights on their own forum. This short paragraph outlines a mechanism for radically distributed storage, distribution, analysis, and reporting on the Diebold files.

As the story unfolded over the next few months, this basic model of peer production of investigation, reportage, analysis, and communication indeed worked. It resulted in the decertification of some of Diebold's systems in California, and contributed to a shift in the requirements of a number of states, which now require voting machines to produce a paper trail for recount purposes. The first analysis of the Diebold system based on the files Harris originally found was performed by a group of computer scientists at the Information Security Institute at Johns Hopkins University and released

as a working paper in late July 2003. The Hopkins Report, or Rubin Report as it was also named after one of its authors, Aviel Rubin, presented deep criticism of the Diebold system and its vulnerabilities on many dimensions. The academic credibility of its authors required a focused response from Diebold. The company published a line-by-line response. Other computer scientists joined in the debate. They showed the limitations and advantages of the Hopkins Report, but also where the Diebold response was adequate and where it provided implicit admission of the presence of a number of the vulnerabilities identified in the report. The report and comments to it sparked two other major reports, commissioned by Maryland in the fall of 2003 and later in January 2004, as part of that state's efforts to decide whether to adopt electronic voting machines. Both studies found a wide range of flaws in the systems they examined and required modifications (see figure 7.2).

Meanwhile, trouble was brewing elsewhere for Diebold. In early August 2003, someone provided *Wired* magazine with a very large cache containing thousands of internal e-mails of Diebold. *Wired* reported that the e-mails were obtained by a hacker, emphasizing this as another example of the laxity of Diebold's security. However, the magazine provided neither an analysis of the e-mails nor access to them. Bev Harris, the activist who had originally found the Diebold materials, on the other hand, received the same cache, and posted the e-mails and memos on her site. Diebold's response was to threaten litigation. Claiming copyright in the e-mails, the company demanded from Harris, her Internet service provider, and a number of other sites where the materials had been posted, that the e-mails be removed. The e-mails were removed from these sites, but the strategy of widely distributed replication of data and its storage in many different topological and organizationally diverse settings made Diebold's efforts ultimately futile. The protagonists from this point on were college students. First, two students at Swarthmore College in Pennsylvania, and quickly students in a number of other universities in the United States, began storing the e-mails and scouring them for evidence of impropriety. In October 2003, Diebold proceeded to write to the universities whose students were hosting the materials. The company invoked provisions of the Digital Millennium Copyright Act that require Web-hosting companies to remove infringing materials when copyright owners notify them of the presence of these materials on their sites. The universities obliged, and required the students to remove the materials from their sites. The students, however, did not disappear quietly into the

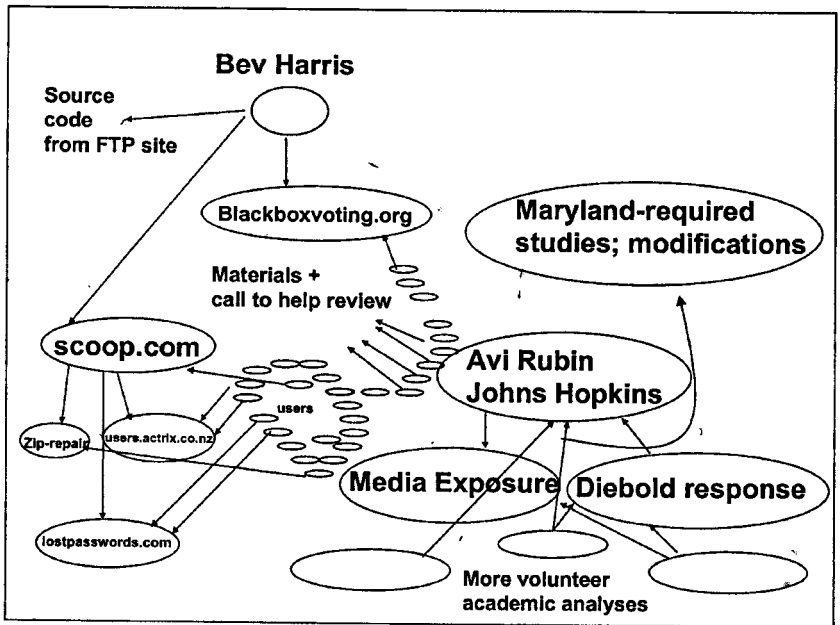


Figure 7.2: Analysis of the Diebold Source Code Materials

night. On October 21, 2003, they launched a multipronged campaign of what they described as “electronic civil disobedience.” First, they kept moving the files from one student to another’s machine, encouraging students around the country to resist the efforts to eliminate the material. Second, they injected the materials into FreeNet, the anticensorship peer-to-peer publication network, and into other peer-to-peer file-sharing systems, like eDonkey and BitTorrent. Third, supported by the Electronic Frontier Foundation, one of the primary civil-rights organizations concerned with Internet freedom, the students brought suit against Diebold, seeking a judicial declaration that their posting of the materials was privileged. They won both the insurgent campaign and the formal one. As a practical matter, the materials remained publicly available throughout this period. As a matter of law, the litigation went badly enough for Diebold that the company issued a letter promising not to sue the students. The court nonetheless awarded the students damages and attorneys’ fees because it found that Diebold had “knowingly and materially misrepresented” that the publication of the e-mail archive was a copyright violation in its letters to the Internet service providers.⁹

Central from the perspective of understanding the dynamics of the networked public sphere is not, however, the court case—it was resolved almost a year later, after most of the important events had already unfolded—but the efficacy of the students' continued persistent publication in the teeth of the cease-and-desist letters and the willingness of the universities to comply. The strategy of replicating the files everywhere made it impracticable to keep the documents from the public eye. And the public eye, in turn, scrutinized. Among the things that began to surface as users read the files were internal e-mails recognizing problems with the voting system, with the security of the FTP site from which Harris had originally obtained the specifications of the voting systems, and e-mail that indicated that the machines implemented in California had been "patched" or updated after their certification. That is, the machines actually being deployed in California were at least somewhat different from the machines that had been tested and certified by the state. This turned out to have been a critical find.

California had a Voting Systems Panel within the office of the secretary of state that reviewed and certified voting machines. On November 3, 2003, two weeks after the students launched their electronic disobedience campaign, the agenda of the panel's meeting was to include a discussion of proposed modifications to one of Diebold's voting systems. Instead of discussing the agenda item, however, one of the panel members made a motion to table the item until the secretary of state had an opportunity to investigate, because "It has come to our attention that some very disconcerting information regarding this item [*sic*] and we are informed that this company, Diebold, may have installed uncertified software in at least one county before it was certified."¹⁰ The source of the information is left unclear in the minutes. A later report in *Wired* cited an unnamed source in the secretary of state's office as saying that somebody within the company had provided this information. The timing and context, however, suggest that it was the revelation and discussion of the e-mail memoranda online that played that role. Two of the members of the public who spoke on the record mention information from within the company. One specifically mentions the information gleaned from company e-mails. In the next committee meeting, on December 16, 2003, one member of the public who was in attendance specifically referred to the e-mails on the Internet, referencing in particular a January e-mail about upgrades and changes to the certified systems. By that December meeting, the independent investigation by the secretary of state had found systematic discrepancies between the systems actually installed

and those tested and certified by the state. The following few months saw more studies, answers, debates, and the eventual decertification of many of the Diebold machines installed in California (see figures 7.3a and 7.3b).

The structure of public inquiry, debate, and collective action exemplified by this story is fundamentally different from the structure of public inquiry and debate in the mass-media-dominated public sphere of the twentieth century. The initial investigation and analysis was done by a committed activist, operating on a low budget and with no financing from a media company. The output of this initial inquiry was not a respectable analysis by a major player in the public debate. It was access to raw materials and initial observations about them, available to start a conversation. Analysis then emerged from a widely distributed process undertaken by Internet users of many different types and abilities. In this case, it included academics studying electronic voting systems, activists, computer systems practitioners, and mobilized students. When the pressure from a well-financed corporation mounted, it was not the prestige and money of a *Washington Post* or a *New York Times* that protected the integrity of the information and its availability for public scrutiny. It was the radically distributed cooperative efforts of students and peer-to-peer network users around the Internet. These efforts were, in turn, nested in other communities of cooperative production—like the free software community that developed some of the applications used to disseminate the e-mails, after Swarthmore removed them from the students' own site. There was no single orchestrating power—neither party nor professional commercial-media outlet. There was instead a series of uncoordinated but mutually reinforcing actions by individuals in different settings and contexts, operating under diverse organizational restrictions and affordances, to expose, analyze, and distribute criticism and evidence for it. The networked public sphere here does not rely on advertising or capturing large audiences to focus its efforts. What became salient for the public agenda and shaped public discussion was what intensely engaged active participants, rather than what kept the moderate attention of large groups of passive viewers. Instead of the lowest-common-denominator focus typical of commercial mass media, each individual and group can—and, indeed, most likely will—focus precisely on what is most intensely interesting to its participants. Instead of iconic representation built on the scarcity of time slots and space on the air or on the page, we see the emergence of a “see for yourself” culture. Access to underlying documents and statements, and to

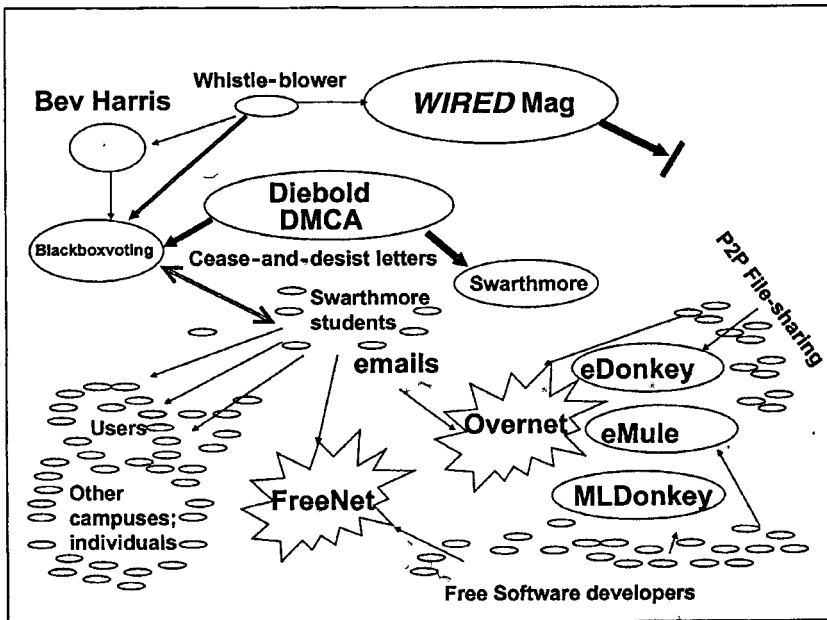


Figure 7.3a: Diebold Internal E-mails, Discovery and Distribution

the direct expression of the opinions of others, becomes a central part of the medium.

CRITIQUES OF THE CLAIMS THAT THE INTERNET HAS DEMOCRATIZING EFFECTS

It is common today to think of the 1990s, out of which came the Supreme Court's opinion in *Reno v. ACLU*, as a time of naïve optimism about the Internet, expressing in political optimism the same enthusiasm that drove the stock market bubble, with the same degree of justifiability. An ideal liberal public sphere did not, in fact, burst into being from the Internet, fully grown like Athena from the forehead of Zeus. The detailed criticisms of the early claims about the democratizing effects of the Internet can be characterized as variants of five basic claims:

1. *Information overload.* A basic problem created when everyone can speak is that there will be too many statements, or too much information. Too

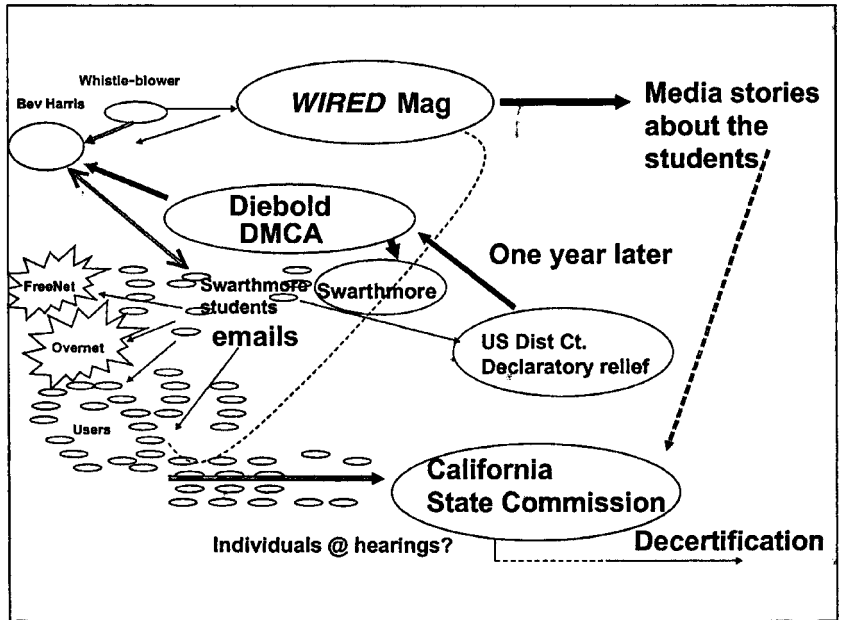


Figure 7.3b: Internal E-mails Translated to Political and Judicial Action

many observations and too many points of view make the problem of sifting through them extremely difficult, leading to an unmanageable din. This overall concern, a variant of the Babel objection, underlies three more specific arguments: that money will end up dominating anyway, that there will be fragmentation of discourse, and that fragmentation of discourse will lead to its polarization.

Money will end up dominating anyway. A point originally raised by Eli Noam is that in this explosively large universe, getting attention will be as difficult as getting your initial message out in the mass-media context, if not more so. The same means that dominated the capacity to speak in the mass-media environment—money—will dominate the capacity to be heard on the Internet, even if it no longer controls the capacity to speak.

Fragmentation of attention and discourse. A point raised most explicitly by Cass Sunstein in *Republic.com* is that the ubiquity of information and the absence of the mass media as condensation points will impoverish public discourse by fragmenting it. There will be no public sphere.

Individuals will view the world through millions of personally customized windows that will offer no common ground for political discourse or action, except among groups of highly similar individuals who customize their windows to see similar things.

Polarization. A descriptively related but analytically distinct critique of Sunstein's was that the fragmentation would lead to polarization. When information and opinions are shared only within groups of like-minded participants, he argued, they tend to reinforce each other's views and beliefs without engaging with alternative views or seeing the concerns and critiques of others. This makes each view more extreme in its own direction and increases the distance between positions taken by opposing camps.

2. *Centralization of the Internet.* A second-generation criticism of the democratizing effects of the Internet is that it turns out, in fact, not to be as egalitarian or distributed as the 1990s conception had suggested. First, there is concentration in the pipelines and basic tools of communications. Second, and more intractable to policy, even in an open network, a high degree of attention is concentrated on a few top sites—a tiny number of sites are read by the vast majority of readers, while many sites are never visited by anyone. In this context, the Internet is replicating the mass-media model, perhaps adding a few channels, but not genuinely changing anything structural.

Note that the concern with information overload is in direct tension with the second-generation concerns. To the extent that the concerns about Internet concentration are correct, they suggest that the information overload is not a deep problem. Sadly, from the perspective of democracy, it turns out that according to the concentration concern, there are few speakers to which most people listen, just as in the mass-media environment. While this means that the supposed benefits of the networked public sphere are illusory, it also means that the information overload concerns about what happens when there is no central set of speakers to whom most people listen are solved in much the same way that the mass-media model deals with the factual diversity of information, opinion, and observations in large societies—by consigning them to public oblivion. The response to both sets of concerns will therefore require combined consideration of a series of questions: To what extent are the claims of concentration correct? How do they solve the information over-

load problem? To what extent does the observed concentration replicate the mass-media model?

3. *Centrality of commercial mass media to the Fourth Estate function.* The importance of the press to the political process is nothing new. It earned the press the nickname "the Fourth Estate" (a reference to the three estates that made up the prerevolutionary French Estates-General, the clergy, nobility, and townsmen), which has been in use for at least a hundred and fifty years. In American free speech theory, the press is often described as fulfilling "the watchdog function," deriving from the notion that the public representatives must be watched over to assure they do the public's business faithfully. In the context of the Internet, the concern, most clearly articulated by Neil Netanel, has been that in the modern complex societies in which we live, commercial mass media are critical for preserving the watchdog function of the media. Big, sophisticated, well-funded government and corporate market actors have enormous resources at their disposal to act as they please and to avoid scrutiny and democratic control. Only similarly big, powerful, independently funded media organizations, whose basic market roles are to observe and criticize other large organizations, can match these established elite organizational actors. Individuals and collections of volunteers talking to each other may be nice, but they cannot seriously replace well-funded, economically and politically powerful media.
4. *Authoritarian countries can use filtering and monitoring to squelch Internet use.* A distinct set of claims and their critiques have to do with the effects of the Internet on authoritarian countries. The critique is leveled at a basic belief supposedly, and perhaps actually, held by some cyber-libertarians, that with enough access to Internet tools freedom will burst out everywhere. The argument is that China, more than any other country, shows that it is possible to allow a population access to the Internet—it is now home to the second-largest national population of Internet users—and still control that use quite substantially.
5. *Digital divide.* While the Internet may increase the circle of participants in the public sphere, access to its tools is skewed in favor of those who already are well-off in society—in terms of wealth, race, and skills. I do not respond to this critique in this chapter. First, in the United States, this is less stark today than it was in the late 1990s. Computers and Internet connections are becoming cheaper and more widely available in public libraries and schools. As they become more central to life, they

seem to be reaching higher penetration rates, and growth rates among underrepresented groups are higher than the growth rate among the highly represented groups. The digital divide with regard to basic access within advanced economies is important as long as it persists, but seems to be a transitional problem. Moreover, it is important to recall that the democratizing effects of the Internet must be compared to democracy in the context of mass media, not in the context of an idealized utopia. Computer literacy and skills, while far from universal, are much more widely distributed than the skills and instruments of mass-media production. Second, I devote chapter 9 to the question of how and why the emergence specifically of nonmarket production provides new avenues for substantial improvements in equality of access to various desiderata that the market distributes unevenly, both within advanced economies and globally, where the maldistribution is much more acute. While the digital divide critique can therefore temper our enthusiasm for how radical the change represented by the networked information economy may be in terms of democracy, the networked information economy is itself an avenue for alleviating maldistribution.

The remainder of this chapter is devoted to responding to these critiques, providing a defense of the claim that the Internet can contribute to a more attractive liberal public sphere. As we work through these objections, we can develop a better understanding of how the networked information economy responds to or overcomes the particular systematic failures of mass media as platforms for the public sphere. Throughout this analysis, it is comparison of the attractiveness of the networked public sphere to that baseline—the mass-media-dominated public sphere—not comparison to a nonexistent ideal public sphere or to the utopia of “everyone a pamphleteer,” that should matter most to our assessment of its democratic promise.

IS THE INTERNET TOO CHAOTIC, TOO CONCENTRATED, OR NEITHER?

The first-generation critique of the claims that the Internet democratizes focused heavily on three variants of the information overload or Babel objection. The basic descriptive proposition that animated the Supreme Court in *Reno v. ACLU* was taken as more or less descriptively accurate: Everyone would be equally able to speak on the Internet. However, this basic obser-

vation was then followed by a descriptive or normative explanation of why this development was a threat to democracy, or at least not much of a boon. The basic problem that is diagnosed by this line of critique is the problem of attention. When everyone can speak, the central point of failure becomes the capacity to be heard—who listens to whom, and how that question is decided. Speaking in a medium that no one will actually hear with any reasonable likelihood may be psychologically satisfying, but it is not a move in a political conversation. Noam's prediction was, therefore, that there would be a reconcentration of attention: money would reemerge in this environment as a major determinant of the capacity to be heard, certainly no less, and perhaps even more so, than it was in the mass-media environment.¹¹ Sunstein's theory was different. He accepted Nicholas Negroponte's prediction that people would be reading "The Daily Me," that is, that each of us would create highly customized windows on the information environment that would be narrowly tailored to our unique combination of interests. From this assumption about how people would be informed, he spun out two distinct but related critiques. The first was that discourse would be fragmented. With no six o'clock news to tell us what is on the public agenda, there would be no public agenda, just a fragmented multiplicity of private agendas that never coalesce into a platform for political discussion. The second was that, in a fragmented discourse, individuals would cluster into groups of self-reinforcing, self-referential discussion groups. These types of groups, he argued from social scientific evidence, tend to render their participants' views more extreme and less amenable to the conversation across political divides necessary to achieve reasoned democratic decisions.

Extensive empirical and theoretical studies of actual use patterns of the Internet over the past five to eight years has given rise to a second-generation critique of the claim that the Internet democratizes. According to this critique, attention is much more concentrated on the Internet than we thought a few years ago: a tiny number of sites are highly linked, the vast majority of "speakers" are not heard, and the democratic potential of the Internet is lost. If correct, these claims suggest that Internet use patterns solve the problem of discourse fragmentation that Sunstein was worried about. Rather than each user reading a customized and completely different "newspaper," the vast majority of users turn out to see the same sites. In a network with a small number of highly visible sites that practically everyone reads, the discourse fragmentation problem is resolved. Because they are seen by most people, the polarization problem too is solved—the highly visible sites are

not small-group interactions with homogeneous viewpoints. While resolving Sunstein's concerns, this pattern is certainly consistent with Noam's prediction that money would have to be paid to reach visibility, effectively replicating the mass-media model. While centralization would resolve the Babel objection, it would do so only at the expense of losing much of the democratic promise of the Net.

Therefore, we now turn to the question: Is the Internet in fact too chaotic or too concentrated to yield a more attractive democratic discourse than the mass media did? I suggest that neither is the case. At the risk of appearing a chimera of Goldilocks and Pangloss, I argue instead that the observed use of the network exhibits an order that is not too concentrated and not too chaotic, but rather, if not "just right," at least structures a networked public sphere more attractive than the mass-media-dominated public sphere.

There are two very distinct types of claims about Internet centralization. The first, and earlier, has the familiar ring of media concentration. It is the simpler of the two, and is tractable to policy. The second, concerned with the emergent patterns of attention and linking on an otherwise open network, is more difficult to explain and intractable to policy. I suggest, however, that it actually stabilizes and structures democratic discourse, providing a better answer to the fears of information overload than either the mass media or any efforts to regulate attention to matters of public concern.

The media-concentration type argument has been central to arguments about the necessity of open access to broadband platforms, made most forcefully over the past few years by Lawrence Lessig. The argument is that the basic instrumentalities of Internet communications are subject to concentrated markets. This market concentration in basic access becomes a potential point of concentration of the power to influence the discourse made possible by access. Eli Noam's recent work provides the most comprehensive study currently available of the degree of market concentration in media industries. It offers a bleak picture.¹² Noam looked at markets in basic infrastructure components of the Internet: Internet backbones, Internet service providers (ISPs), broadband providers, portals, search engines, browser software, media player software, and Internet telephony. Aggregating across all these sectors, he found that the Internet sector defined in terms of these components was, throughout most of the period from 1984 to 2002, concentrated according to traditional antitrust measures. Between 1992 and 1998, however, this sector was "highly concentrated" by the Justice Department's measure of market concentration for antitrust purposes. Moreover, the power

of the top ten firms in each of these markets, and in aggregate for firms that had large market segments in a number of these markets, shows that an ever-smaller number of firms were capturing about 25 percent of the revenues in the Internet sector. A cruder, but consistent finding is the FCC's, showing that 96 percent of homes and small offices get their broadband access either from their incumbent cable operator or their incumbent local telephone carrier.¹³ It is important to recognize that these findings are suggesting potential points of failure for the networked information economy. They are not a critique of the democratic potential of the networked public sphere, but rather show us how we could fail to develop it by following the wrong policies.

The risk of concentration in broadband access services is that a small number of firms, sufficiently small to have economic power in the antitrust sense, will control the markets for the basic instrumentalities of Internet communications. Recall, however, that the low cost of computers and the open-ended architecture of the Internet protocol itself are the core enabling facts that have allowed us to transition from the mass-media model to the networked information model. As long as these basic instrumentalities are open and neutral as among uses, and are relatively cheap, the basic economics of nonmarket production described in part I should not change. Under competitive conditions, as technology makes computation and communications cheaper, a well-functioning market should ensure that outcome. Under oligopolistic conditions, however, there is a threat that the network will become too expensive to be neutral as among market and nonmarket production. If basic upstream network connections, server space, and up-to-date reading and writing utilities become so expensive that one needs to adopt a commercial model to sustain them, then the basic economic characteristic that typifies the networked information economy—the relatively large role of nonproprietary, nonmarket production—will have been reversed. However, the risk is not focused solely or even primarily on explicit pricing. One of the primary remaining scarce resources in the networked environment is user time and attention. As chapter 5 explained, owners of communications facilities can extract value from their users in ways that are more subtle than increasing price. In particular, they can make some sites and statements easier to reach and see—more prominently displayed on the screen, faster to load—and sell that relative ease to those who are willing to pay.¹⁴ In that environment, nonmarket sites are systematically disadvantaged irrespective of the quality of their content.

The critique of concentration in this form therefore does not undermine the claim that the networked information economy, if permitted to flourish, will improve the democratic public sphere. It underscores the threat of excessive monopoly in infrastructure to the sustainability of the networked public sphere. The combination of observations regarding market concentration and an understanding of the importance of a networked public sphere to democratic societies suggests that a policy intervention is possible and desirable. Chapter 11 explains why the relevant intervention is to permit substantial segments of the core common infrastructure—the basic physical transport layer of wireless or fiber and the software and standards that run communications—to be produced and provisioned by users and managed as a commons.

ON POWER LAW DISTRIBUTIONS, NETWORK TOPOLOGY, AND BEING HEARD

A much more intractable challenge to the claim that the networked information economy will democratize the public sphere emerges from observations of a set of phenomena that characterize the Internet, the Web, the blogosphere, and, indeed, most growing networks. In order to extract information out of the universe of statements and communications made possible by the Internet, users are freely adopting practices that lead to the emergence of a new hierarchy. Rather than succumb to the “information overload” problem, users are solving it by congregating in a small number of sites. This conclusion is based on a new but growing literature on the likelihood that a Web page will be linked to by others. The distribution of that probability turns out to be highly skew. That is, there is a tiny probability that any given Web site will be linked to by a huge number of people, and a very large probability that for a given Web site only one other site, or even no site, will link to it. This fact is true of large numbers of very different networks described in physics, biology, and social science, as well as in communications networks. If true in this pure form about Web usage, this phenomenon presents a serious theoretical and empirical challenge to the claim that Internet communications of the sorts we have seen here meaningfully decentralize democratic discourse. It is not a problem that is tractable to policy. We cannot as a practical matter force people to read different things than what they choose to read; nor should we wish to. If users avoid information overload by focusing on a small subset of sites in an otherwise

open network that allows them to read more or less whatever they want and whatever anyone has written, policy interventions aimed to force a different pattern would be hard to justify from the perspective of liberal democratic theory.

The sustained study of the distribution of links on the Internet and the Web is relatively new—only a few years old. There is significant theoretical work in a field of mathematics called graph theory, or network topology, on power law distributions in networks, on skew distributions that are not pure power law, and on the mathematically related small-worlds phenomenon in networks. The basic intuition is that, if indeed a tiny minority of sites gets a large number of links, and the vast majority gets few or no links, it will be very difficult to be seen unless you are on the highly visible site. Attention patterns make the open network replicate mass media. While explaining this literature over the next few pages, I show that what is in fact emerging is very different from, and more attractive than, the mass-media-dominated public sphere.

While the Internet, the Web, and the blogosphere are indeed exhibiting much greater order than the freewheeling, “everyone a pamphleteer” image would suggest, this structure does not replicate a mass-media model. We are seeing a newly shaped information environment, where indeed few are read by many, but clusters of moderately read sites provide platforms for vastly greater numbers of speakers than were heard in the mass-media environment. Filtering, accreditation, synthesis, and salience are created through a system of peer review by information affinity groups, topical or interest based. These groups filter the observations and opinions of an enormous range of people, and transmit those that pass local peer review to broader groups and ultimately to the polity more broadly, without recourse to market-based points of control over the information flow. Intense interest and engagement by small groups that share common concerns, rather than lowest-common-denominator interest in wide groups that are largely alienated from each other, is what draws attention to statements and makes them more visible. This makes the emerging networked public sphere more responsive to intensely held concerns of a much wider swath of the population than the mass media were capable of seeing, and creates a communications process that is more resistant to corruption by money.

In what way, first, is attention concentrated on the Net? We are used to seeing probability distributions that describe social phenomena following a Gaussian distribution: where the mean and the median are the same and the

probabilities fall off symmetrically as we describe events that are farther from the median. This is the famous Bell Curve. Some phenomena, however, observed initially in Pareto's work on income distribution and Zipf's on the probability of the use of English words in text and in city populations, exhibit completely different probability distributions. These distributions have very long "tails"—that is, they are characterized by a very small number of very high-yield events (like the number of words that have an enormously high probability of appearing in a randomly chosen sentence, like "the" or "to") and a very large number of events that have a very low probability of appearing (like the probability that the word "probability" or "blogosphere" will appear in a randomly chosen sentence). To grasp intuitively how un-intuitive such distributions are to us, we could think of radio humorist Garrison Keillor's description of the fictitious Lake Wobegon, where "all the children are above average." That statement is amusing because we assume intelligence follows a normal distribution. If intelligence were distributed according to a power law, most children there would actually be below average—the median is well below the mean in such distributions (see figure 7.4). Later work by Herbert Simon in the 1950s, and by Derek de Solla Price in the 1960s, on cumulative advantage in scientific citations¹⁵ presaged an emergence at the end of the 1990s of intense interest in power law characterizations of degree distributions, or the number of connections any point in a network has to other points, in many kinds of networks—from networks of neurons and axons, to social networks and communications and information networks.

The Internet and the World Wide Web offered a testable setting, where large-scale investigation could be done automatically by studying link structure (who is linked-in to and by whom, who links out and to whom, how these are related, and so on), and where the practical applications of better understanding were easily articulated—such as the design of better search engines. In 1999, Albert-László Barabási and Reka Albert published a paper in *Science* showing that a variety of networked phenomena have a predictable topology: The distribution of links into and out of nodes on the network follows a power law. There is a very low probability that any vertex, or node, in the network will be very highly connected to many others, and a very large probability that a very large number of nodes will be connected only very loosely, or perhaps not at all. Intuitively, a lot of Web sites link to information that is located on Yahoo!, while very few link to any randomly selected individual's Web site. Barabási and Albert hypothesized a mechanism

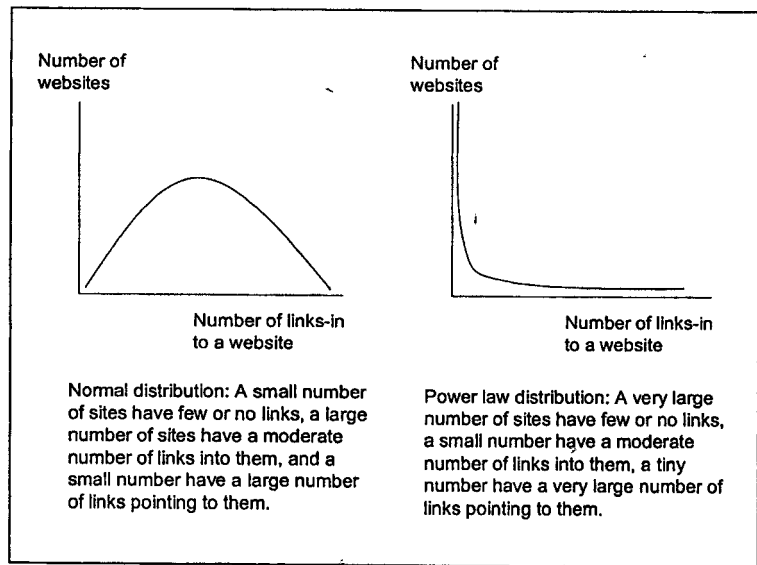


Figure 7.4: Illustration of How Normal Distribution and Power Law Distribution Would Differ in Describing How Many Web Sites Have Few or Many Links Pointing at Them

for this distribution to evolve, which they called “preferential attachment.” That is, new nodes prefer to attach to already well-attached nodes. Any network that grows through the addition of new nodes, and in which nodes preferentially attach to nodes that are already well attached, will eventually exhibit this distribution.¹⁶ In other words, the rich get richer. At the same time, two computer scientists, Lada Adamic and Bernardo Huberman, published a study in *Nature* that identified the presence of power law distributions in the number of Web pages in a given site. They hypothesized not that new nodes preferentially attach to old ones, but that each site has an intrinsically different growth rate, and that new sites are formed at an exponential rate.¹⁷ The intrinsically different growth rates could be interpreted as quality, interest, or perhaps investment of money in site development and marketing. They showed that on these assumptions, a power law distribution would emerge. Since the publication of these articles we have seen an explosion of theoretical and empirical literature on graph theory, or the structure and growth of networks, and particularly on link structure in the World Wide Web. It has consistently shown that the number of links into and out of Web sites follows power laws and that the exponent (the expo-

nential factor that determines that the drop-off between the most linked-to site and the second most linked-to site, and the third, and so on, will be so dramatically rapid, and how rapid it is) for inlinks is roughly 2.1 and for outlinks 2.7.

If one assumes that most people read things by either following links, or by using a search engine, like Google, that heavily relies on counting inlinks to rank its results, then it is likely that the number of visitors to a Web page, and more recently, the number of readers of blogs, will follow a similarly highly skew distribution. The implication for democracy that comes most immediately to mind is dismal. While, as the Supreme Court noted with enthusiasm, on the Internet everyone can be a pamphleteer or have their own soapbox, the Internet does not, in fact, allow individuals to be heard in ways that are substantially more effective than standing on a soapbox in a city square. Many Web pages and blogs will simply go unread, and will not contribute to a more engaged polity. This argument was most clearly made in Barabási's popularization of his field, *Linked*: "The most intriguing result of our Web-mapping project was the *complete* absence of democracy, fairness, and egalitarian values on the Web. We learned that the topology of the Web prevents us from seeing anything but a mere handful of the billion documents out there."¹⁸

The stories offered in this chapter and throughout this book present a puzzle for this interpretation of the power law distribution of links in the network as re-creating a concentrated medium. The success of Nick Davis's site, BoycottSBG, would be a genuine fluke. The probability that such a site could be established on a Monday, and by Friday of the same week would have had three hundred thousand unique visitors and would have orchestrated a successful campaign, is so small as to be negligible. The probability that a completely different site, StopSinclair.org, of equally network-obscure origins, would be established on the very same day and also successfully catch the attention of enough readers to collect 150,000 signatures on a petition to protest Sinclair's broadcast, rather than wallowing undetected in the mass of self-published angry commentary, is practically insignificant. And yet, intuitively, it seems unsurprising that a large population of individuals who are politically mobilized on the same side of the political map and share a political goal in the public sphere—using a network that makes it trivially simple to set up new points of information and coordination, tell each other about them, and reach and use them from anywhere—would, in fact, inform each other and gather to participate in a political demonstration. We saw

that the boycott technique that Davis had designed his Web site to facilitate was discussed on TalkingPoints—a site near the top of the power law distribution of political blogs—but that it was a proposal by an anonymous individual who claimed to know what makes local affiliates tick, not of TalkingPoints author Josh Marshall. By midweek, after initially stoking the fires of support for Davis's boycott, Marshall had stepped back, and Davis's site became the clearing point for reports, tactical conversations, and mobilization. Davis not only was visible, but rather than being drowned out by the high-powered transmitter, TalkingPoints, his relationship with the high-visibility site was part of his success. This story alone cannot, of course, "refute" the power law distribution of network links, nor is it offered as a refutation. It does, however, provide a context for looking more closely at the emerging understanding of the topology of the Web, and how it relates to the fears of concentration of the Internet, and the problems of information overload, discourse fragmentation, and the degree to which money will come to dominate such an unstructured and wide-open environment. It suggests a more complex story than simply "the rich get richer" and "you might speak, but no one will hear you." In this case, the topology of the network allowed rapid emergence of a position, its filtering and synthesis, and its rise to salience. Network topology helped facilitate all these components of the public sphere, rather than undermined them. We can go back to the mathematical and computer science literature to begin to see why.

Within two months of the publication of Barabási and Albert's article, Adamic and Huberman had published a letter arguing that, if Barabási and Albert were right about preferential attachment, then older sites should systematically be among those that are at the high end of the distribution, while new ones will wallow in obscurity. The older sites are already attached, so newer sites would preferentially attach to the older sites. This, in turn, would make them even more attractive when a new crop of Web sites emerged and had to decide which sites to link to. In fact, however, Adamic and Huberman showed that there is no such empirical correlation among Web sites. They argued that their mechanism—that nodes have intrinsic growth rates that are different—better describes the data. In their response, Barabási and Albert showed that on their data set, the older nodes are actually more connected in a way that follows a power law, but only on average—that is to say, the average number of connections of a class of older nodes related to the average number of links to a younger class of nodes follows a power law. This argued that their basic model was sound, but

required that they modify their equations to include something similar to what Huberman and Adamic had proposed—an intrinsic growth factor for each node, as well as the preferential connection of new nodes to established nodes.¹⁹ This modification is important because it means that not every new node is doomed to be unread relative to the old ones, only that on average they are much less likely to be read. It makes room for rapidly growing new nodes, but does not theorize what might determine the rate of growth. It is possible, for example, that money could determine growth rates: In order to be seen, new sites or statements would have to spend money to gain visibility and salience. As the BoycottSBG and Diebold stories suggest, however, as does the Lott story described later in this chapter, there are other ways of achieving immediate salience. In the case of BoycottSBG, it was providing a solution that resonated with the political beliefs of many people and was useful to them for their expression and mobilization. Moreover, the continued presence of preferential attachment suggests that noncommercial Web sites that are already highly connected because of the time they were introduced (like the Electronic Frontier Foundation), because of their internal attractiveness to large communities (like Slashdot), or because of their salience to the immediate interests of users (like BoycottSBG), will have persistent visibility even in the face of large infusions of money by commercial sites.

Developments in network topology theory and its relationship to the structure of the empirically mapped real Internet offer a map of the networked information environment that is indeed quite different from the naïve model of “everyone a pamphleteer.” To the limited extent that these findings have been interpreted for political meaning, they have been seen as a disappointment—the real world, as it turns out, does not measure up to anything like that utopia. However, that is the wrong baseline. There never has been a complex, large modern democracy in which everyone could speak and be heard by everyone else. The correct baseline is the one-way structure of the commercial mass media. The normatively relevant descriptive questions are whether the networked public sphere provides broader intake, participatory filtering, and relatively incorruptible platforms for creating public salience. I suggest that it does. Four characteristics of network topology structure the Web and the blogosphere in an ordered, but nonetheless meaningfully participatory form. First, at a microlevel, sites cluster—in particular, topically and interest-related sites link much more heavily to each other than to other sites. Second, at a macrolevel, the Web and the blogosphere have

giant, strongly connected cores—"areas" where 20–30 percent of all sites are highly and redundantly interlinked; that is, tens or hundreds of millions of sites, rather than ten, fifty, or even five hundred television stations. That pattern repeats itself in smaller subclusters as well. Third, as the clusters get small enough, the obscurity of sites participating in the cluster diminishes, while the visibility of the superstars remains high, forming a filtering and transmission backbone for universal intake and local filtering. Fourth and finally, the Web exhibits "small-world" phenomena, making most Web sites reachable through shallow paths from most other Web sites. I will explain each of these below, as well as how they interact to form a reasonably attractive image of the networked public sphere.

First, links are not smoothly distributed throughout the network. Sites cluster into densely linked "regions" or communities of interest. Computer scientists have looked at clustering from the perspective of what topical or other correlated characteristics describe these relatively high-density interconnected regions of nodes. What they found was perhaps entirely predictable from an intuitive perspective of the network users, but important as we try to understand the structure of information flow on the Web. Web sites cluster into topical and social/organizational clusters. Early work done in the IBM Almaden Research Center on how link structure could be used as a search technique showed that by mapping densely interlinked sites without looking at content, one could find communities of interest that identify very fine-grained topical connections, such as Australian fire brigades or Turkish students in the United States.²⁰ A later study out of the NEC Research Institute more formally defined the interlinking that would identify a "community" as one in which the nodes were more densely connected to each other than they were to nodes outside the cluster by some amount. The study also showed that topically connected sites meet this definition. For instance, sites related to molecular biology clustered with each other—in the sense of being more interlinked with each other than with off-topic sites—as did sites about physics and black holes.²¹ Lada Adamic and Natalie Glance recently showed that liberal political blogs and conservative political blogs densely interlink with each other, mostly pointing within each political leaning but with about 15 percent of links posted by the most visible sites also linking across the political divide.²² Physicists analyze clustering as the property of transitivity in networks: the increased probability that if node A is connected to node B, and node B is connected to node C, that node A also will be connected to node C, forming a triangle. Newman has shown that

the clustering coefficient of a network that exhibits power law distribution of connections or degrees—that is, its tendency to cluster—is related to the exponent of the distribution. At low exponents, below 2.333, the clustering coefficient becomes high. This explains analytically the empirically observed high level of clustering on the Web, whose exponent for inlinks has been empirically shown to be 2.1.²³

Second, at a macrolevel and in smaller subclusters, the power law distribution does not resolve into everyone being connected in a mass-media model relationship to a small number of major “backbone” sites. As early as 1999, Broder and others showed that a very large number of sites occupy what has been called a giant, strongly connected core.²⁴ That is, nodes within this core are heavily linked and interlinked, with multiple redundant paths among them. Empirically, as of 2001, this structure was comprised of about 28 percent of nodes. At the same time, about 22 percent of nodes had links into the core, but were not linked to from it—these may have been new sites, or relatively lower-interest sites. The same proportion of sites was linked-to from the core, but did not link back to it—these might have been ultimate depositories of documents, or internal organizational sites. Finally, roughly the same proportion of sites occupied “tendrils” or “tubes” that cannot reach, or be reached from, the core. Tendrils can be reached from the group of sites that link into the strongly connected core or can reach into the group that can be connected to from the core. Tubes connect the inlinking sites to the outlinked sites without going through the core. About 10 percent of sites are entirely isolated. This structure has been called a “bow tie”—with a large core and equally sized in- and outflows to and from that core (see figure 7.5).

One way of interpreting this structure as counterdemocratic is to say: This means that half of all Web sites are not reachable from the other half—the “IN,” “tendrils,” and disconnected portions cannot be reached from any of the sites in SCC and OUT. This is indeed disappointing from the “everyone a pamphleteer” perspective. On the other hand, one could say that half of all Web pages, the SCC and OUT components, *are* reachable from IN and SCC. That is, hundreds of millions of pages are reachable from hundreds of millions of potential entry points. This represents a very different intake function and freedom to speak in a way that is potentially accessible to others than a five-hundred-channel, mass-media model. More significant yet, Dill and others showed that the bow tie structure appears not only at the level of the Web as a whole, but repeats itself within clusters. That is, the Web




Figure 7.5: Bow Tie Structure of the Web

appears to show characteristics of self-similarity, up to a point—links within clusters also follow a power law distribution and cluster, and have a bow tie structure of similar proportions to that of the overall Web. Tying the two points about clustering and the presence of a strongly connected core, Dill and his coauthors showed that what they called “thematically unified clusters,” such as geographically or content-related groupings of Web sites, themselves exhibit these strongly connected cores that provided a thematically defined navigational backbone to the Web. It is not that one or two major sites were connected to by all thematically related sites; rather, as at the network level, on the order of 25–30 percent were highly interlinked, and another 25 percent were reachable from within the strongly connected core.²⁵ Moreover, when the data was pared down to treat only the home page, rather than each Web page within a single site as a distinct “node” (that is, everything that came under *www.foo.com* was treated as one node, as opposed to the usual method where *www.foo.com*, *www.foo.com/nonsuch*, and *www.foo.com/somethingelse* are each treated as a separate node), fully 82 percent of the nodes were in the strongly connected core, and an additional 13 percent were reachable from the SCC as the OUT group.

Third, another finding of Web topology and critical adjustment to the

basic Barabási and Albert model is that when the topically or organizationally related clusters become small enough—on the order of hundreds or even low thousands of Web pages—they no longer follow a pure power law distribution. Instead, they follow a distribution that still has a very long tail—these smaller clusters still have a few genuine “superstars”—but the body of the distribution is substantially more moderate: beyond the few superstars, the shape of the link distribution looks a little more like a normal distribution. Instead of continuing to drop off exponentially, many sites exhibit a moderate degree of connectivity. Figure 7.6 illustrates how a hypothetical distribution of this sort would differ both from the normal and power law distributions illustrated in figure 7.4. David Pennock and others, in their paper describing these empirical findings, hypothesized a uniform component added to the purely exponential original Barabási and Albert model. This uniform component could be random (as they modeled it), but might also stand for quality of materials, or level of interest in the site by participants in the smaller cluster. At large numbers of nodes, the exponent dominates the uniform component, accounting for the pure power law distribution when looking at the Web as a whole, or even at broadly defined topics. In smaller clusters of sites, however, the uniform component begins to exert a stronger pull on the distribution. The exponent keeps the long tail intact, but the uniform component accounts for a much more moderate body. Many sites will have dozens, or even hundreds of links. The Pennock paper looked at sites whose number was reduced by looking only at sites of certain organizations—universities or public companies. Chakrabarti and others later confirmed this finding for topical clusters as well. That is, when they looked at small clusters of topically related sites, the distribution of links still has a long tail for a small number of highly connected sites in every topic, but the body of the distribution diverges from a power law distribution, and represents a substantial proportion of sites that are moderately linked.²⁶ Even more specifically, Daniel Drezner and Henry Farrell reported that the Pennock modification better describes distribution of links specifically to and among political blogs.²⁷

These findings are critical to the interpretation of the distribution of links as it relates to human attention and communication. There is a big difference between a situation where no one is looking at any of the sites on the low end of the distribution, because everyone is looking only at the superstars, and a situation where dozens or hundreds of sites at the low end are looking at each other, as well as at the superstars. The former leaves all but the very




Figure 7.6: Illustration of a Skew Distribution That Does Not Follow a Power Law

few languishing in obscurity, with no one to look at them. The latter, as explained in more detail below, offers a mechanism for topically related and interest-based clusters to form a peer-reviewed system of filtering, accreditation, and salience generation. It gives the long tail on the low end of the distribution heft (and quite a bit of wag).

The fourth and last piece of mapping the network as a platform for the public sphere is called the “small-worlds effect.” Based on Stanley Milgram’s sociological experiment and on mathematical models later proposed by Duncan Watts and Steven Strogatz, both theoretical and empirical work has shown that the number of links that must be traversed from any point in the network to any other point is relatively small.²⁸ Fairly shallow “walks”—that is, clicking through three or four layers of links—allow a user to cover a large portion of the Web.

What is true of the Web as a whole turns out to be true of the blogosphere as well, and even of the specifically political blogosphere. Early 2003 saw increasing conversations in the blogosphere about the emergence of an “A-list,” a number of highly visible blogs that were beginning to seem more like mass media than like blogs. In two blog-based studies, Clay Shirky and then Jason Kottke published widely read explanations of how the blogo-

sphere was simply exhibiting the power law characteristics common on the Web.²⁹ The emergence in 2003 of discussions of this sort in the blogosphere is, it turns out, hardly surprising. In a time-sensitive study also published in 2003, Kumar and others provided an analysis of the network topology of the blogosphere. They found that it was very similar to that of the Web as a whole—both at the macro- and microlevels. Interestingly, they found that the strongly connected core only developed after a certain threshold, in terms of total number of nodes, had been reached, and that it began to develop extensively only in 2001, reached about 20 percent of all blogs in 2002, and continued to grow rapidly. They also showed that what they called the “community” structure—the degree of clustering or mutual pointing within groups—was high, an order of magnitude more than a random graph with a similar power law exponent would have generated. Moreover, the degree to which a cluster is active or inactive, highly connected or not, changes over time. In addition to time-insensitive superstars, there are also flare-ups of connectivity for sites depending on the activity and relevance of their community of interest. This latter observation is consistent with what we saw happen for BoycottSBG.com. Kumar and his collaborators explained these phenomena by the not-too-surprising claim that bloggers link to each other based on topicality—that is, their judgment of the quality and relevance of the materials—not only on the basis of how well connected they are already.³⁰

This body of literature on network topology suggests a model for how order has emerged on the Internet, the World Wide Web, and the blogosphere. The networked public sphere allows hundreds of millions of people to publish whatever and whenever they please without disintegrating into an unusable cacophony, as the first-generation critics argued, and it filters and focuses attention without re-creating the highly concentrated model of the mass media that concerned the second-generation critique. We now know that the network at all its various layers follows a degree of order, where some sites are vastly more visible than most. This order is loose enough, however, and exhibits a sufficient number of redundant paths from an enormous number of sites to another enormous number, that the effect is fundamentally different from the small number of commercial professional editors of the mass media.

Individuals and individual organizations cluster around topical, organizational, or other common features. At a sufficiently fine-grained degree of clustering, a substantial proportion of the clustered sites are moderately con-

nected, and each can therefore be a point of intake that will effectively transmit observations or opinions within and among the users of that topical or interest-based cluster. Because even in small clusters the distribution of links still has a long tail, these smaller clusters still include high-visibility nodes. These relatively high-visibility nodes can serve as points of transfer to larger clusters, acting as an attention backbone that transmits information among clusters. Subclusters within a general category—such as liberal and conservative blogs clustering within the broader cluster of political blogs—are also interlinked, though less densely than within-cluster connectivity. The higher level or larger clusters again exhibit a similar feature, where higher visibility nodes can serve as clearinghouses and connectivity points among clusters and across the Web. These are all highly connected with redundant links within a giant, strongly connected core—comprising more than a quarter of the nodes in any given level of cluster. The small-worlds phenomenon means that individual users who travel a small number of different links from similar starting points within a cluster cover large portions of the Web and can find diverse sites. By then linking to them on their own Web sites, or giving them to others by e-mail or blog post, sites provide multiple redundant paths open to many users to and from most statements on the Web. High-visibility nodes amplify and focus on given statements, and in this regard, have greater power in the information environment they occupy. However, there is sufficient redundancy of paths through high-visibility nodes that no single node or small collection of nodes can control the flow of information in the core and around the Web. This is true both at the level of the cluster and at the level of the Web as a whole.

The result is an ordered system of intake, filtering, and synthesis that can in theory emerge in networks generally, and empirically has been shown to have emerged on the Web. It does not depend on single points of control. It avoids the generation of a din through which no voice can be heard, as the fears of fragmentation predicted. And, while money may be useful in achieving visibility, the structure of the Web means that money is neither necessary nor sufficient to grab attention—because the networked information economy, unlike its industrial predecessor, does not offer simple points of dissemination and control for purchasing assured attention. What the network topology literature allows us to do, then, is to offer a richer, more detailed, and empirically supported picture of how the network can be a platform for the public sphere that is structured in a fundamentally different way than the mass-media model. The problem is approached

through a self-organizing principle, beginning with communities of interest on smallish scales, practices of mutual pointing, and the fact that, with freedom to choose what to see and who to link to, with some codependence among the choices of individuals as to whom to link, highly connected points emerge even at small scales, and continue to be replicated with ever-larger visibility as the clusters grow. Without forming or requiring a formal hierarchy, and without creating single points of control, each cluster generates a set of sites that offer points of initial filtering, in ways that are still congruent with the judgments of participants in the highly connected small cluster. The process is replicated at larger and more general clusters, to the point where positions that have been synthesized “locally” and “regionally” can reach Web-wide visibility and salience. It turns out that we are not intellectual lemmings. We do not use the freedom that the network has made possible to plunge into the abyss of incoherent babble. Instead, through iterative processes of cooperative filtering and “transmission” through the high visibility nodes, the low-end thin tail turns out to be a peer-produced filter and transmission medium for a vastly larger number of speakers than was imaginable in the mass-media model.

The effects of the topology of the network are reinforced by the cultural forms of linking, e-mail lists, and the writable Web. The network topology literature treats every page or site as a node. The emergence of the writable Web, however, allows each node to itself become a cluster of users and posters who, collectively, gain salience as a node. Slashdot is “a node” in the network as a whole, one that is highly linked and visible. Slashdot itself, however, is a highly distributed system for peer production of observations and opinions about matters that people who care about information technology and communications ought to care about. Some of the most visible blogs, like the dailyKos, are cooperative blogs with a number of authors. More important, the major blogs receive input—through posts or e-mails—from their users. Recall, for example, that the original discussion of a Sinclair boycott that would focus on local advertisers arrived on TalkingPoints through an e-mail comment from a reader. Talkingpoints regularly solicits and incorporates input from and research by its users. The cultural practice of writing to highly visible blogs with far greater ease than writing a letter to the editor and with looser constraints on what gets posted makes these nodes themselves platforms for the expression, filtering, and synthesis of observations and opinions. Moreover, as Drezner and Farrell have shown, blogs have developed cultural practices of mutual citation—when one blog-

ger finds a source by reading another, the practice is to link to the original blog, not only directly to the underlying source. Jack Balkin has argued that the culture of linking more generally and the "see for yourself" culture also significantly militate against fragmentation of discourse, because users link to materials they are commenting on, even in disagreement.

Our understanding of the emerging structure of the networked information environment, then, provides the basis for a response to the family of criticisms of the first generation claims that the Internet democratizes. Recall that these criticisms, rooted in the problem of information overload, or the Babel objection, revolved around three claims. The first claim was that the Internet would result in a fragmentation of public discourse. The clustering of topically related sites, such as politically oriented sites, and of communities of interest, the emergence of high-visibility sites that the majority of sites link to, and the practices of mutual linking show quantitatively and qualitatively what Internet users likely experience intuitively. While there is enormous diversity on the Internet, there are also mechanisms and practices that generate a common set of themes, concerns, and public knowledge around which a public sphere can emerge. Any given site is likely to be within a very small number of clicks away from a site that is visible from a very large number of other sites, and these form a backbone of common materials, observations, and concerns. All the findings of power law distribution of linking, clustering, and the presence of a strongly connected core, as well as the linking culture and "see for yourself," oppose the fragmentation prediction. Users self-organize to filter the universe of information that is generated in the network. This self-organization includes a number of highly salient sites that provide a core of common social and cultural experiences and knowledge that can provide the basis for a common public sphere, rather than a fragmented one.

The second claim was that fragmentation would cause polarization. Because like-minded people would talk only to each other, they would tend to amplify their differences and adopt more extreme versions of their positions. Given that the evidence demonstrates there is no fragmentation, in the sense of a lack of a common discourse, it would be surprising to find higher polarization because of the Internet. Moreover, as Balkin argued, the fact that the Internet allows widely dispersed people with extreme views to find each other and talk is not a failure for the liberal public sphere, though it may present new challenges for the liberal state in constraining extreme action. Only polarization of discourse in society as a whole can properly be

considered a challenge to the attractiveness of the networked public sphere. However, the practices of linking, "see for yourself," or quotation of the position one is criticizing, and the widespread practice of examining and criticizing the assumptions and assertions of one's interlocutors actually point the other way, militating against polarization. A potential counterargument, however, was created by the most extensive recent study of the political blogosphere. In that study, Adamic and Glance showed that only about 10 percent of the links on any randomly selected political blog linked to a site across the ideological divide. The number increased for the "A-list" political blogs, which linked across the political divide about 15 percent of the time. The picture that emerges is one of distinct "liberal" and "conservative" spheres of conversation, with very dense links within, and more sparse links between them. On one interpretation, then, although there are salient sites that provide a common subject matter for discourse, actual conversations occur in distinct and separate spheres,—exactly the kind of setting that Sunstein argued would lead to polarization. Two of the study's findings, however, suggest a different interpretation. The first was that there was still a substantial amount of cross-divide linking. One out of every six or seven links in the top sites on each side of the divide linked to the other side in roughly equal proportions (although conservatives tended to link slightly more overall—both internally and across the divide). The second was, that in an effort to see whether the more closely interlinked conservative sites therefore showed greater convergence "on message," Adamic and Glance found that greater interlinking did not correlate with less diversity in external (outside of the blogosphere) reference points.³¹ Together, these findings suggest a different interpretation. Each cluster of more or less like-minded blogs tended to read each other and quote each other much more than they did the other side. This operated not so much as an echo chamber as a forum for working out of observations and interpretations internally, among like-minded people. Many of these initial statements or inquiries die because the community finds them uninteresting or fruitless. Some reach greater salience, and are distributed through the high-visibility sites throughout the community of interest. Issues that in this form reached political salience became topics of conversation and commentary across the divide. This is certainly consistent with both the BoycottSBG and Diebold stories, where we saw a significant early working out of strategies and observations before the criticism reached genuine political salience. There would have been no point for opponents to link to and criticize early ideas kicked around within the com-

munity, like opposing Sinclair station renewal applications. Only after a few days, when the boycott was crystallizing, would opponents have reason to point out the boycott effort and discuss it. This interpretation also well characterizes the way in which the Trent Lott story described later in this chapter began percolating on the liberal side of the blogosphere, but then migrated over to the center-right.

The third claim was that money would reemerge as the primary source of power brokerage because of the difficulty of getting attention on the Net. Descriptively, it shares a prediction with the second-generation claims: Namely, that the Internet will centralize discourse. It differs in the mechanism of concentration: it will not be the result of an emergent property of large-scale networks, but rather of an old, tried-and-true way of capturing the political arena—money. But the peer-production model of filtering and discussion suggests that the networked public sphere will be substantially less corruptible by money. In the interpretation that I propose, filtering for the network as a whole is done as a form of nested peer-review decisions, beginning with the speaker's closest information affinity group. Consistent with what we have been seeing in more structured peer-production projects like *Wikipedia*, Slashdot, or free software, communities of interest use clustering and mutual pointing to peer produce the basic filtering mechanism necessary for the public sphere to be effective and avoid being drowned in the din of the crowd. The nested structure of the Web, whereby subclusters form relatively dense higher-level clusters, which then again combine into even higher-level clusters, and in each case, have a number of high-end salient sites, allows for the statements that pass these filters to become globally salient in the relevant public sphere. This structure, which describes the analytic and empirical work on the Web as a whole, fits remarkably well as a description of the dynamics we saw in looking more closely at the success of the boycott on Sinclair, as well as the successful campaign to investigate and challenge Diebold's voting machines.

The peer-produced structure of the attention backbone suggests that money is neither necessary nor sufficient to attract attention in the networked public sphere (although nothing suggests that money has become irrelevant to political attention given the continued importance of mass media). It renders less surprising Howard Dean's strong campaign for the Democratic presidential primaries in 2003 and the much more stable success of MoveOn.org since the late 1990s. These suggest that attention on the network has more to do with mobilizing the judgments, links, and cooperation

of large bodies of small-scale contributors than with applying large sums of money. There is no obvious broadcast station that one can buy in order to assure salience. There are, of course, the highly visible sites, and they do offer a mechanism of getting your message to large numbers of people. However, the degree of engaged readership, interlinking, and clustering suggests that, in fact, being exposed to a certain message in one or a small number of highly visible places accounts for only a small part of the range of "reading" that gets done. More significantly, it suggests that reading, as opposed to having a conversation, is only part of what people do in the networked environment. In the networked public sphere, receiving information or getting out a finished message are only parts, and not necessarily the most important parts, of democratic discourse. The central desideratum of a political campaign that is rooted in the Internet is the capacity to engage users to the point that they become effective participants in a conversation and an effort; one that they have a genuine stake in and that is linked to a larger, society-wide debate. This engagement is not easily purchased, nor is it captured by the concept of a well-educated public that receives all the information it needs to be an informed citizenry. Instead, it is precisely the varied modes of participation in small-, medium-, and large-scale conversations, with varied but sustained degrees of efficacy, that make the public sphere of the networked environment different, and more attractive, than was the mass-media-based public sphere.

The networked public sphere is not only more resistant to control by money, but it is also less susceptible to the lowest-common-denominator orientation that the pursuit of money often leads mass media to adopt. Because communication in peer-produced media starts from an intrinsic motivation—writing or commenting about what one cares about—it begins with the opposite of lowest common denominator. It begins with what irks you, the contributing peer, individually, the most. This is, in the political world, analogous to Eric Raymond's claim that every free or open-source software project begins with programmers with an itch to scratch—something directly relevant to their lives and needs that they want to fix. The networked information economy, which makes it possible for individuals alone and in cooperation with others to scour the universe of politically relevant events, to point to them, and to comment and argue about them, follows a similar logic. This is why one freelance writer with lefty leanings, Russ Kick, is able to maintain a Web site, *The Memory Hole*, with documents that he gets by filing Freedom of Information Act requests. In April

2004, Kick was the first to obtain the U.S. military's photographs of the coffins of personnel killed in Iraq being flown home. No mainstream news organization had done so, but many published the photographs almost immediately after Kick had obtained them. Like free software, like Davis and the bloggers who participated in the debates over the Sinclair boycott, or the students who published the Diebold e-mails, the decision of what to publish does not start from a manager's or editor's judgment of what would be relevant and interesting to many people without being overly upsetting to too many others. It starts with the question: What do I care about most now?

To conclude, we need to consider the attractiveness of the networked public sphere not from the perspective of the mid-1990s utopianism, but from the perspective of how it compares to the actual media that have dominated the public sphere in all modern democracies. The networked public sphere provides an effective nonmarket alternative for intake, filtering, and synthesis outside the market-based mass media. This nonmarket alternative can attenuate the influence over the public sphere that can be achieved through control over, or purchase of control over, the mass media. It offers a substantially broader capture basin for intake of observations and opinions generated by anyone with a stake in the polity, anywhere. It appears to have developed a structure that allows for this enormous capture basin to be filtered, synthesized, and made part of a polity-wide discourse. This nested structure of clusters of communities of interest, typified by steadily increasing visibility of superstar nodes, allows for both the filtering and salience to climb up the hierarchy of clusters, but offers sufficient redundant paths and interlinking to avoid the creation of a small set of points of control where power can be either directly exercised or bought.

There is, in this story, an enormous degree of contingency and factual specificity. That is, my claims on behalf of the networked information economy as a platform for the public sphere are not based on general claims about human nature, the meaning of liberal discourse, context-independent efficiency, or the benevolent nature of the technology we happen to have stumbled across at the end of the twentieth century. They are instead based on, and depend on the continued accuracy of, a description of the economics of fabrication of computers and network connections, and a description of the dynamics of linking in a network of connected nodes. As such, my claim is not that the Internet inherently liberates. I do not claim that commons-based production of information, knowledge, and culture will win out by

some irresistible progressive force. That is what makes the study of the political economy of information, knowledge, and culture in the networked environment directly relevant to policy. The literature on network topology suggests that, as long as there are widely distributed capabilities to publish, link, and advise others about what to read and link to, networks enable intrinsic processes that allow substantial ordering of the information. The pattern of information flow in such a network is more resistant to the application of control or influence than was the mass-media model. But things can change. Google could become so powerful on the desktop, in the e-mail utility, and on the Web, that it will effectively become a supernode that will indeed raise the prospect of a reemergence of a mass-media model. Then the politics of search engines, as Lucas Introna and Helen Nissenbaum called it, become central. The zeal to curb peer-to-peer file sharing of movies and music could lead to a substantial redesign of computing equipment and networks, to a degree that would make it harder for end users to exchange information of their own making. Understanding what we will lose if such changes indeed warp the topology of the network, and through it the basic structure of the networked public sphere, is precisely the object of this book as a whole. For now, though, let us say that the networked information economy as it has developed to this date has a capacity to take in, filter, and synthesize observations and opinions from a population that is orders of magnitude larger than the population that was capable of being captured by the mass media. It has done so without re-creating identifiable and reliable points of control and manipulation that would replicate the core limitation of the mass-media model of the public sphere—its susceptibility to the exertion of control by its regulators, owners, or those who pay them.

WHO WILL PLAY THE WATCHDOG FUNCTION?

A distinct critique leveled at the networked public sphere as a platform for democratic politics is the concern for who will fill the role of watchdog. Neil Netanel made this argument most clearly. His concern was that, perhaps freedom of expression for all is a good thing, and perhaps we could even overcome information overflow problems, but we live in a complex world with powerful actors. Government and corporate power is large, and individuals, no matter how good their tools, cannot be a serious alternative to a well-funded, independent press that can pay investigative reporters, defend lawsuits, and generally act like the *New York Times* and the *Washington Post*

when they published the Pentagon Papers in the teeth of the Nixon administration's resistance, providing some of the most damning evidence against the planning and continued prosecution of the war in Vietnam. Netanel is cognizant of the tensions between the need to capture large audiences and sell advertising, on the one hand, and the role of watchdog, on the other. He nonetheless emphasizes that the networked public sphere cannot investigate as deeply or create the public salience that the mass media can. These limitations make commercial mass media, for all their limitations, necessary for a liberal public sphere.

This diagnosis of the potential of the networked public sphere underrepresents its productive capacity. The Diebold story provides in narrative form a detailed response to each of the concerns. The problem of voting machines has all the characteristics of an important, hard subject. It stirs deep fears that democracy is being stolen, and is therefore highly unsettling. It involves a difficult set of technical judgments about the functioning of voting machines. It required exposure and analysis of corporate-owned materials in the teeth of litigation threats and efforts to suppress and discredit the criticism. At each juncture in the process, the participants in the critique turned iteratively to peer production and radically distributed methods of investigation, analysis, distribution, and resistance to suppression: the initial observations of the whistle-blower or the hacker; the materials made available on a "see for yourself" and "come analyze this and share your insights" model; the distribution by students; and the fallback option when their server was shut down of replication around the network. At each stage, a peer-production solution was interposed in place of where a well-funded, high-end mass-media outlet would have traditionally applied funding in expectation of sales of copy. And it was only after the networked public sphere developed the analysis and debate that the mass media caught on, and then only gingerly.

The Diebold case was not an aberration, but merely a particularly-rich case study of a much broader phenomenon, most extensively described in Dan Gilmore's *We the Media*. The basic production modalities that typify the networked information economy are now being applied to the problem of producing politically relevant information. In 2005, the most visible example of application of the networked information economy—both in its peer-production dimension and more generally by combining a wide range of nonproprietary production models—to the watchdog function of the media is the political blogosphere. The founding myth of the blogosphere's

journalistic potency was built on the back of then Senate majority leader Trent Lott. In 2002, Lott had the indiscretion of saying, at the one-hundredth-birthday party of Republican Senator Strom Thurmond, that if Thurmond had won his Dixiecrat presidential campaign, "we wouldn't have had all these problems over all these years." Thurmond had run on a segregationist campaign, splitting from the Democratic Party in opposition to Harry Truman's early civil rights efforts, as the post-World War II winds began blowing toward the eventual demise of formal, legal racial segregation in the United States. Few positions are taken to be more self-evident in the national public morality of early twenty-first-century America than that formal, state-imposed, racial discrimination is an abomination. And yet, the first few days after the birthday party at which Lott made his statement saw almost no reporting on the statement. ABC News and the *Washington Post* made small mention of it, but most media outlets reported merely on a congenial salute and farewell celebration of the Senate's oldest and longest-serving member. Things were different in the blogosphere. At first liberal blogs, and within three days conservative bloggers as well, began to excavate past racist statements by Lott, and to beat the drums calling for his censure or removal as Senate leader. Within about a week, the story surfaced in the mainstream media, became a major embarrassment, and led to Lott's resignation as Senate majority leader about a week later. A careful case study of this event leaves it unclear why the mainstream media initially ignored the story.³² It may have been that the largely social event drew the wrong sort of reporters. It may have been that reporters and editors who depend on major Washington, D.C., players were reluctant to challenge Lott. Perhaps they thought it rude to emphasize this indiscretion, or too upsetting to us all to think of just how close to the surface thoughts that we deem abominable can lurk. There is little disagreement that the day after the party, the story was picked up and discussed by Marshall on TalkingPoints, as well as by another liberal blogger, Atrios, who apparently got it from a post on Slate's "Chatterbox," which picked it up from ABC News's own *The Note*, a news summary made available on the television network's Web site. While the mass media largely ignored the story, and the two or three mainstream reporters who tried to write about it were getting little traction, bloggers were collecting more stories about prior instances where Lott's actions tended to suggest support for racist causes. Marshall, for example, found that Lott had filed a 1981 amicus curiae brief in support of Bob Jones University's effort to retain its tax-exempt status. The U.S. government had rescinded

that status because the university practiced racial discrimination—such as prohibiting interracial dating. By Monday of the following week, four days after the remarks, conservative bloggers like Glenn Reynolds on Instapundit, Andrew Sullivan, and others were calling for Lott's resignation. It is possible that, absent the blogosphere, the story would still have flared up. There were two or so mainstream reporters still looking into the story. Jesse Jackson had come out within four days of the comment and said Lott should resign as majority leader. Eventually, when the mass media did enter the fray, its coverage clearly dominated the public agenda and its reporters uncovered materials that helped speed Lott's exit. However, given the short news cycle, the lack of initial interest by the media, and the large time lag between the event itself and when the media actually took the subject up, it seems likely that without the intervention of the blogosphere, the story would have died. What happened instead is that the cluster of political blogs—starting on the Left but then moving across the Left-Right divide—took up the subject, investigated, wrote opinions, collected links and public interest, and eventually captured enough attention to make the comments a matter of public importance. Free from the need to appear neutral and not to offend readers, and free from the need to keep close working relationships with news subjects, bloggers were able to identify something that grated on their sensibilities, talk about it, dig deeper, and eventually generate a substantial intervention into the public sphere. That intervention still had to pass through the mass media, for we still live in a communications environment heavily based on those media. However, the new source of insight, debate, and eventual condensation of effective public opinion came from within the networked information environment.

The point is not to respond to the argument with a litany of anecdotes. The point is that the argument about the commercial media's role as watchdog turns out to be a familiar argument—it is the same argument that was made about software and supercomputers, encyclopedias and immersive entertainment scripts. The answer, too, is by now familiar. Just as the World Wide Web can offer a platform for the emergence of an enormous and effective almanac, just as free software can produce excellent software and peer production can produce a good encyclopedia, so too can peer production produce the public watchdog function. In doing so, clearly the unorganized collection of Internet users lacks some of the basic tools of the mass media: dedicated full-time reporters; contacts with politicians who need media to survive, and therefore cannot always afford to stonewall questions; or

public visibility and credibility to back their assertions. However, network-based peer production also avoids the inherent conflicts between investigative reporting and the bottom line—its cost, its risk of litigation, its risk of withdrawal of advertising from alienated corporate subjects, and its risk of alienating readers. Building on the wide variation and diversity of knowledge, time, availability, insight, and experience, as well as the vast communications and information resources on hand for almost anyone in advanced economies, we are seeing that the watchdog function too is being peer produced in the networked information economy.

Note that while my focus in this chapter has been mostly the organization of public discourse, both the Sinclair and the Diebold case studies also identify characteristics of distributed political action. We see collective action emerging from the convergence of independent individual actions, with no hierarchical control like that of a political party or an organized campaign. There may be some coordination and condensation points—like [BoycottSBG.com](#) or [blackboxvoting.org](#). Like other integration platforms in peer-production systems, these condensation points provide a critical function. They do not, however, control the process. One manifestation of distributed coordination for political action is something Howard Rheingold has called “smart mobs”—large collections of individuals who are able to coordinate real-world action through widely distributed information and communications technology. He tells of the “People Power II” revolution in Manila in 2001, where demonstrations to oust then president Estrada were coordinated spontaneously through extensive text messaging.³³ Few images in the early twenty-first century can convey this phenomenon more vividly than the demonstrations around the world on February 15, 2003. Between six and ten million protesters were reported to have gone to the streets of major cities in about sixty countries in opposition to the American-led invasion of Iraq. There had been no major media campaign leading up to the demonstrations—though there was much media attention to them later. There had been no organizing committee. Instead, there was a network of roughly concordant actions, none controlling the other, all loosely discussing what ought to be done and when. [MoveOn.org](#) in the United States provides an example of a coordination platform for a network of politically mobilized activities. It builds on e-mail and Web-based media to communicate opportunities for political action to those likely to be willing and able to take it. Radically distributed, network-based solutions to the problems of political mobilization rely on the same characteristics as networked information production

more generally: extensive communications leading to concordant and cooperative patterns of behavior without the introduction of hierarchy or the interposition of payment.

USING NETWORKED COMMUNICATION TO WORK AROUND AUTHORITARIAN CONTROL

The Internet and the networked public sphere offer a different set of potential benefits, and suffer a different set of threats, as a platform for liberation in authoritarian countries. State-controlled mass-media models are highly conducive to authoritarian control. Because they usually rely on a small number of technical and organizational points of control, mass media offer a relatively easy target for capture and control by governments. Successful control of such universally visible media then becomes an important tool of information manipulation, which, in turn, eases the problem of controlling the population. Not surprisingly, capture of the national television and radio stations is invariably an early target of coups and revolutions. The highly distributed networked architecture of the Internet makes it harder to control communications in this way.

The case of Radio B92 in Yugoslavia offers an example. B92 was founded in 1989, as an independent radio station. Over the course of the 1990s, it developed a significant independent newsroom broadcast over the station itself, and syndicated through thirty affiliated independent stations. B92 was banned twice after the NATO bombing of Belgrade, in an effort by the Milosevic regime to control information about the war. In each case, however, the station continued to produce programming, and distributed it over the Internet from a server based in Amsterdam. The point is a simple one. Shutting down a broadcast station is simple. There is one transmitter with one antenna, and police can find and hold it. It is much harder to shut down all connections from all reporters to a server and from the server back into the country wherever a computer exists.

This is not to say that the Internet will of necessity in the long term lead all authoritarian regimes to collapse. One option open to such regimes is simply to resist Internet use. In 2003, Burma, or Myanmar, had 28,000 Internet users out of a population of more than 42 million, or one in fifteen hundred, as compared, for example, to 6 million out of 65 million in neighboring Thailand, or roughly one in eleven. Most countries are not, however, willing to forgo the benefits of connectivity to maintain their control. Iran's

population of 69 million includes 4.3 million Internet users, while China has about 80 million users, second only to the United States in absolute terms, out of a population of 1.3 billion. That is, both China and Iran have a density of Internet users of about one in sixteen.³⁴ Burma's negligible level of Internet availability is a compound effect of low gross domestic product (GDP) per capita and government policies. Some countries with similar GDP levels still have levels of Internet users in the population that are two orders of magnitude higher: Cameroon (1 Internet user for every 27 residents), Moldova (1 in 30), and Mongolia (1 in 55). Even very large poor countries have several times more users per population than Myanmar: like Pakistan (1 in 100), Mauritania (1 in 300), and Bangladesh (1 in 580). Lawrence Solum and Minn Chung outline how Myanmar achieves its high degree of control and low degree of use.³⁵ Myanmar has only one Internet service provider (ISP), owned by the government. The government must authorize anyone who wants to use the Internet or create a Web page within the country. Some of the licensees, like foreign businesses, are apparently permitted and enabled only to send e-mail, while using the Web is limited to security officials who monitor it. With this level of draconian regulation, Myanmar can avoid the liberating effects of the Internet altogether, at the cost of losing all its economic benefits. Few regimes are willing to pay that price.

Introducing Internet communications into a society does not, however, immediately and automatically mean that an open, liberal public sphere emerges. The Internet is technically harder to control than mass media. It increases the cost and decreases the efficacy of information control. However, a regime willing and able to spend enough money and engineering power, and to limit its population's access to the Internet sufficiently, can have substantial success in controlling the flow of information into and out of its country. Solum and Chung describe in detail one of the most extensive and successful of these efforts, the one that has been conducted by China—home to the second-largest population of Internet users in the world, whose policies controlled use of the Internet by two out of every fifteen Internet users in the world in 2003. In China, the government holds a monopoly over all Internet connections going into and out of the country. It either provides or licenses the four national backbones that carry traffic throughout China and connect it to the global network. ISPs that hang off these backbones are licensed, and must provide information about the location and workings of their facilities, as well as comply with a code of conduct. In-

dividual users must register and provide information about their machines, and the many Internet cafes are required to install filtering software that will filter out subversive sites. There have been crackdowns on Internet cafes to enforce these requirements. This set of regulations has replicated one aspect of the mass-medium model for the Internet—it has created a potential point of concentration or centralization of information flow that would make it easier to control Internet use. The highly distributed production capabilities of the networked information economy, however, as opposed merely to the distributed carriage capability of the Internet, mean that more must be done at this bottleneck to squelch the flow of information and opinion than would have to be done with mass media. That “more” in China has consisted of an effort to employ automatic filters—some at the level of the cybercafe or the local ISP, some at the level of the national backbone networks. The variability of these loci and their effects is reflected in partial efficacy and variable performance for these mechanisms. The most extensive study of the efficacy of these strategies for controlling information flows over the Internet to China was conducted by Jonathan Zittrain and Ben Edelman. From servers within China, they sampled about two hundred thousand Web sites and found that about fifty thousand were unavailable at least once, and close to nineteen thousand were unavailable on two distinct occasions. The blocking patterns seemed to follow mass-media logic—BBC News was consistently unavailable, as CNN and other major news sites often were; the U.S. court system official site was unavailable. However, Web sites that provided similar information—like those that offered access to all court cases but were outside the official system—were available. The core Web sites of human rights organizations or of Taiwan and Tibet-related organizations were blocked, and about sixty of the top one hundred results for “Tibet” on Google were blocked. What is also apparent from their study, however, and confirmed by Amnesty International’s reports on Internet censorship in China, is that while censorship is significant, it is only partially effective.³⁶ The Amnesty report noted that Chinese users were able to use a variety of techniques to avoid the filtering, such as the use of proxy servers, but even Zittrain and Edelman, apparently testing for filtering as experienced by unsophisticated or compliant Internet users in China, could access many sites that would, on their face, seem potentially destabilizing.

This level of censorship may indeed be effective enough for a government negotiating economic and trade expansion with political stability and control. It suggests, however, limits of the ability of even a highly dedicated

government to control the capacity of Internet communications to route around censorship and to make it much easier for determined users to find information they care about, and to disseminate their own information to others. Iran's experience, with a similar level of Internet penetration, emphasizes the difficulty of maintaining control of Internet publication.³⁷ Iran's network emerged from 1993 onward from the university system, quite rapidly complemented by commercial ISPs. Because deployment and use of the Internet preceded its regulation by the government, its architecture is less amenable to centralized filtering and control than China's. Internet access through university accounts and cybercafes appears to be substantial, and until the past three or four years, had operated free of the crackdowns and prison terms suffered by opposition print publications and reporters. The conservative branches of the regime seem to have taken a greater interest in suppressing Internet communications since the publication of imprisoned Ayatollah Montazeri's critique of the foundations of the Islamic state on the Web in December 2000. While the original Web site, *montazeri.com*, seems to have been eliminated, the site persists as *montazeri.ws*, using a Western Samoan domain name, as do a number of other Iranian publications. There are now dozens of chat rooms, blogs, and Web sites, and e-mail also seems to be playing an increasing role in the education and organization of an opposition. While the conservative branches of the Iranian state have been clamping down on these forms, and some bloggers and Web site operators have found themselves subject to the same mistreatment as journalists, the efficacy of these efforts to shut down opposition seems to be limited and uneven.

Media other than static Web sites present substantially deeper problems for regimes like those of China and Iran. Scanning the text of e-mail messages of millions of users who can encrypt their communications with widely available tools creates a much more complex problem. Ephemeral media like chat rooms and writable Web tools allow the content of an Internet communication or Web site to be changed easily and dynamically, so that blocking sites becomes harder, while coordinating moves to new sites to route around blocking becomes easier. At one degree of complexity deeper, the widely distributed architecture of the Net also allows users to build censorship-resistant networks by pooling their own resources. The pioneering example of this approach is Freenet, initially developed in 1999–2000 by Ian Clarke, an Irish programmer fresh out of a degree in computer science and artificial intelligence at Edinburgh University. Now a broader free-software project, Freenet

is a peer-to-peer application specifically designed to be censorship resistant. Unlike the more famous peer-to-peer network developed at the time—Napster—Freenet was not intended to store music files on the hard drives of users. Instead, it stores bits and pieces of publications, and then uses sophisticated algorithms to deliver the documents to whoever seeks them, in encrypted form. This design trades off easy availability for a series of security measures that prevent even the owners of the hard drives on which the data resides—or government agents that search their computers—from knowing what is on their hard drive or from controlling it. As a practical matter, if someone in a country that prohibits certain content but enables Internet connections wants to publish content—say, a Web site or blog—safely, they can inject it into the Freenet system. The content will be encrypted and divided into little bits and pieces that are stored in many different hard drives of participants in the network. No single computer will have all the information, and shutting down any given computer will not make the information unavailable. It will continue to be accessible to anyone running the Freenet client. Freenet indeed appears to be used in China, although the precise scope is hard to determine, as the network is intended to mask the identity and location of both readers and publishers in this system. The point to focus on is not the specifics of Freenet, but the feasibility of constructing user-based censorship-resistant storage and retrieval systems that would be practically impossible for a national censorship system to identify and block subversive content.

To conclude, in authoritarian countries, the introduction of Internet communications makes it harder and more costly for governments to control the public sphere. If these governments are willing to forgo the benefits of Internet connectivity, they can avoid this problem. If they are not, they find themselves with less control over the public sphere. There are, obviously, other means of more direct repression. However, control over the mass media was, throughout most of the twentieth century, a core tool of repressive governments. It allowed them to manipulate what the masses of their populations knew and believed, and thus limited the portion of the population that the government needed to physically repress to a small and often geographically localized group. The efficacy of these techniques of repression is blunted by adoption of the Internet and the emergence of a networked information economy. Low-cost communications, distributed technical and organizational structure, and ubiquitous presence of dynamic authorship

tools make control over the public sphere difficult, and practically never perfect.

TOWARD A NETWORKED PUBLIC SPHERE

The first generation of statements that the Internet democratizes was correct but imprecise. The Internet does restructure public discourse in ways that give individuals a greater say in their governance than the mass media made possible. The Internet does provide avenues of discourse around the bottlenecks of older media, whether these are held by authoritarian governments or by media owners. But the mechanisms for this change are more complex than those articulated in the past. And these more complex mechanisms respond to the basic critiques that have been raised against the notion that the Internet enhances democracy.

Part of what has changed with the Internet is technical infrastructure. Network communications do not offer themselves up as easily for single points of control as did the mass media. While it is possible for authoritarian regimes to try to retain bottlenecks in the Internet, the cost is higher and the efficacy lower than in mass-media-dominated systems. While this does not mean that introduction of the Internet will automatically result in global democratization, it does make the work of authoritarian regimes harder. In liberal democracies, the primary effect of the Internet runs through the emergence of the networked information economy. We are seeing the emergence to much greater significance of nonmarket, individual, and cooperative peer-production efforts to produce universal intake of observations and opinions about the state of the world and what might and ought to be done about it. We are seeing the emergence of filtering, accreditation, and synthesis mechanisms as part of network behavior. These rely on clustering of communities of interest and association and highlighting of certain sites, but offer tremendous redundancy of paths for expression and accreditation. These practices leave no single point of failure for discourse: no single point where observations can be squelched or attention commanded—by fiat or with the application of money. Because of these emerging systems, the networked information economy is solving the information overload and discourse fragmentation concerns without reintroducing the distortions of the mass-media model. Peer production, both long-term and organized, as in the case of Slashdot, and ad hoc and dynamically formed, as in the case of blogging or

the Sinclair or Diebold cases, is providing some of the most important functionalities of the media. These efforts provide a watchdog, a source of salient observations regarding matters of public concern, and a platform for discussing the alternatives open to a polity.

In the networked information environment, everyone is free to observe, report, question, and debate, not only in principle, but in actual capability. They can do this, if not through their own widely read blog, then through a cycle of mailing lists, collective Web-based media like Slashdot, comments on blogs, or even merely through e-mails to friends who, in turn, have meaningful visibility in a smallish-scale cluster of sites or lists. We are witnessing a fundamental change in how individuals can interact with their democracy and experience their role as citizens. Ideal citizens need not be seen purely as trying to inform themselves about what others have found, so that they can vote intelligently. They need not be limited to reading the opinions of opinion makers and judging them in private conversations. They are no longer constrained to occupy the role of mere readers, viewers, and listeners. They can be, instead, participants in a conversation. Practices that begin to take advantage of these new capabilities shift the locus of content creation from the few professional journalists trolling society for issues and observations, to the people who make up that society. They begin to free the public agenda setting from dependence on the judgments of managers, whose job it is to assure that the maximum number of readers, viewers, and listeners are sold in the market for eyeballs. The agenda thus can be rooted in the life and experience of individual participants in society—in their observations, experiences, and obsessions. The network allows all citizens to change their relationship to the public sphere. They no longer need be consumers and passive spectators. They can become creators and primary subjects. It is in this sense that the Internet democratizes.