

UseNet and Democratization

Bryan Ford

Handout (Draft) — September 7, 2026

1 Why Start Here

Why begin a course on technologies for democratic society with a network most students have never used, and which faded before most of them were born? Because UseNet is the first case, and in some ways the purest, of a pattern this course will return to again and again.

UseNet was a worldwide discussion system built before the Web, running on infrastructure no one owned, open to anyone who could reach it. To many who used it, it looked like something democracy had long promised and rarely delivered: a place where anyone could speak to everyone, with no editor, publisher, or government in between. That belief was not naive. UseNet genuinely did things no earlier medium could. But it also strained in ways its builders never anticipated, and eventually it broke. Which of its promises held, which failed, and what made the difference, is the work of this chapter.

That arc—promise, then strain, then reckoning—is not unique to UseNet. Nearly every technology in this course, from blockchain governance to social media to artificial intelligence, has been greeted as a democratizing force, and each has faced its own version of the same test. UseNet offers a complete case, with decades of hindsight, on which to learn to read the pattern before turning to stories still unfolding.

2 A Platform for Publishing

UseNet began in 1979 as a link between Unix machines at Duke University and the University of North Carolina, built by three graduate students, Tom Truscott, Jim Ellis, and Steve Bellovin [13]. The idea was simple. Each site periodically dialed its neighbors over ordinary telephone lines and exchanged whatever messages were new. A message posted anywhere would, hop by hop, eventually reach everywhere. There was no central server and no one in charge; a site joined by finding one existing site willing to exchange news with it.

Messages were sorted into *newsgroups*, named topics such as `comp.lang.c` or `rec.arts.sf`, organized into hierarchies: `comp.*` for computing, `sci.*` for science, `soc.*` for social issues, `rec.*` for recreation, and later the deliberately unruly `alt.*`. Anyone could read any group and post to any group.

Consider what that meant. Every earlier mass medium had a gatekeeper. To reach a large audience through print, radio, or television you needed an editor, a publisher, or a broadcast license, and those gatekeepers decided whose voice was worth carrying. UseNet removed them. For the first time in history, everyone with access was not merely a reader but a publisher. A student at a small

college could post to `sci.physics` and be read, on equal terms, by a Nobel laureate; the laureate could reply, and so could anyone else. Whenever you liked, on whatever you liked, you could address the world.

And the platform was reliable in a way no gatekept medium could be. News flowed along many overlapping paths, so no single failure or refusal could cut a site off; robustness grew with the number of alternative routes between any two points, and no government or company could suppress a post short of severing every route at once. The design also made UseNet slow and thin: early links ran at 300 bits per second, and a post could take days to propagate. But it worked. From that modest start UseNet grew, by the early 1990s, into thousands of newsgroups on an enormous diversity of topics, read by millions of people at universities, companies, and, increasingly, homes.

This was the first of UseNet's two roles, and on its own terms it succeeded. The democratization of publication, to everyone at once, with no one able to stop it, was real, and it was one source of the enormous excitement UseNet generated. Keep that success in view. It is the foundation on which a larger expectation was built.

3 A Platform for Self-Organization

A broadcast medium lets people speak. UseNet quickly became something more: a place where people talked *back*, argued, worked problems through together, and sometimes tried to decide things. Its earliest users were programmers and system administrators, and the earliest groups were about the problems such people care about: which compiler to use, how to fix a device driver, whether a proposed protocol was sound. Questions were asked and answered by strangers, disagreements ran for weeks, and consensus, when it came, came from the thread itself rather than from any authority. Over time the same pattern spread far beyond computing, to science, politics, religion, hobbies, and support groups for people who had never before found anyone else with their problem.

Here a second and larger expectation grew out of the first. UseNet had already proven itself as an open, resilient publication medium that no one controlled. If it could do that, why could it not also become a genuine tool for collective self-governance, a way for communities to deliberate and decide without any central authority at all? The absence of government came to be seen not as a gap waiting to be filled but as a feature. The community would govern itself.

The evidence that people took this seriously is that they tried. Virginia Shea's *Netiquette* [15] is the clearest artifact: a codification of norms, from "remember the human" to "respect other people's time and bandwidth," offered not as rules enforced by anyone but as a shared understanding newcomers were expected to absorb. Groups wrote and maintained FAQs. A formal proposal-and-vote procedure for creating newsgroups emerged. Hauben and Hauben [8] coined "netizen" to name the resulting sense of civic membership in a place with no state. That any of this existed at all, self-generated and unenforced, was an achievement in its own right; how far it went, and where it broke, occupies much of the rest of this chapter.

It also raises the question this course keeps asking. A medium in which anyone can speak is not yet a democracy; a democracy is a way of deciding. Did UseNet's self-organization amount to real democratic self-government, or only to something that resembled it? To answer that we need a standard for what "real" would mean, which is where Dahl comes in. First, though, the hope itself.

4 The Promise

Before turning to standards, pause on the mood. The excitement UseNet produced is distilled in John Perry Barlow’s 1996 “Declaration of the Independence of Cyberspace” [2], addressed to the “Governments of the Industrial World, you weary giants of flesh and steel”: “You have no sovereignty where we gather. . . . We are creating a world that all may enter without privilege or prejudice accorded by race, economic power, military force, or station of birth.” Barlow meant the internet broadly, but the sentiment was UseNet’s native idealism made explicit. Read it twice. Was Barlow describing a place that actually existed, UseNet as it was by 1996, or an aspiration that no medium before or since has fully achieved? Which sentences are observation, and which are hope?

The deeper roots of that hope lie in an essay written a decade before UseNet existed. In 1968 Licklider and Taylor [12] imagined a future of computer-mediated communication, and much of what they imagined UseNet became. So put the question to them directly: what would Licklider make of UseNet? Answering well means going back to specific claims in the essay and testing each against what you now know.

- The essay opens by predicting that “in a few years, men will be able to communicate more effectively through a machine than face to face.” Did UseNet bear this out? More effectively for what, and compared with which alternatives?
- Licklider and Taylor foresaw online communities “not of common location, but of common interest.” UseNet’s newsgroups are the obvious realization. What did such communities gain by dropping geography, and what did they lose?
- The essay closes with a warning: the network’s effect on society depends on whether “to be on line” turns out to be “a privilege or a right.” Which was it, for UseNet, over its first fifteen years? Does the answer tell against the prediction, or only against its timing?
- They predicted, only half joking, that the work of building and maintaining the network would make unemployment “disappear from the face of the earth forever.” How has that aged, and does UseNet offer any evidence either way?
- Finally, consider their playful aside about OLIVER, an automated agent that would read, filter, and answer a person’s correspondence on their behalf: something like ChatGPT, fifty years early. UseNet had no OLIVER. Would it have wanted one, and for what? Hold that question; it returns at the end of this chapter and again later in this course.

Do not settle these now. Write down your provisional answers, and check them against the rest of the chapter.

5 What Would Dahl Make of UseNet?

Robert Dahl [4] proposed five criteria that a process must satisfy to count as fully democratic. He wrote them for governments, but they apply to any collective arrangement, and asking them of UseNet is revealing precisely because UseNet was not a government at all. So: what would Dahl make of UseNet? For each criterion, assemble the evidence on both sides before deciding. You should find that none of the five yields a clean verdict.

Effective participation. Members must have adequate and equal opportunity to express their views before a decision is reached. On its face, UseNet is the strongest case for this criterion any medium had yet offered: anyone with access could post to any group, at any length, and be carried to every site in the world. But Dahl’s wording is careful. Is the opportunity to *post* the same as the opportunity to be *heard*? And “before a decision is reached”: participation in which decision, exactly?

Voting equality at the decisive stage. When a decision is made, every member’s vote must count equally. Start by asking whether UseNet reached a decisive stage at all. Creating a newsgroup was eventually put to a vote, cast by email, counted by volunteers, and open to anyone who noticed; most of what mattered—which posts were read, which norms held, who was ignored—was never decided by anyone. Does a process that mostly avoids decisions satisfy this criterion, evade it, or expose a gap in Dahl’s framework?

Enlightened understanding. Members must have adequate opportunity to learn what the alternatives are and what they would mean. UseNet had FAQs, archives, and experts answering questions for free. It also had every posted opinion sitting beside every other, with nothing to mark which was informed. Does a medium meet this criterion by making information available, or must it also make it findable? And what happens to your answer as the volume grows?

Control of the agenda. Members must be able to decide what gets decided. Anyone could start a thread or propose a group. Yet someone chose the hierarchy names, and someone decided which groups each site would carry. Who set UseNet’s agenda, formally and in practice, and were they the same people? The next section supplies a test case.

Inclusiveness. Everyone subject to the process should be included in it. Ask who could reach UseNet in 1985, in which countries, and in what proportions by gender, language, and wealth. Then ask Dahl’s harder question: subject to *what* process? If UseNet decided nothing binding, who was “subject” to it, and does the criterion even apply? And if your answer to the first question is “a narrow slice of humanity,” is that a passing condition of a young technology or a structural feature of “open” systems?

Keep your five provisional answers. The rest of this chapter is evidence, and Section 11 returns to the scorecard.

6 Governance Without Government: The Great Renaming

For most of its first decade, UseNet’s structure was tended by a loose group of administrators at the sites that carried the most traffic, half-jokingly known as the “Backbone Cabal.” Two changes then reshaped the network in quick succession. In 1986, the Network News Transfer Protocol (NNTP) let news travel over the Internet rather than by dial-up, greatly expanding its reach. In 1987 the Cabal carried out the Great Renaming, reorganizing the ad hoc groups that had accumulated into seven top-level hierarchies (**comp**, **sci**, **soc**, **rec**, **talk**, **news**, **misc**) with a formal proposal-and-vote procedure for creating new groups.

The renaming was practical housekeeping, but it was also an exercise of power. Contentious topics were shunted into **talk.***, which many sites then declined to carry, and the Cabal decided which proposed groups were acceptable. When proposals for groups on sex and on recreational drugs were refused, Brian Reid and John Gilmore responded not by appealing but by routing around. They created the **alt.*** hierarchy, where anyone could create a group with no vote at all (**alt.sex**, **alt.drugs**, and, Reid added, for completeness, **alt.rock-n-roll**), and persuaded

enough sites to carry it that it became permanent.

Look at this through Dahl’s *control of the agenda*. Here was a genuine contest over who decides what the community may discuss, with no constitution, no court, and no authority able to settle it. It was resolved not by procedure but by architecture: because UseNet was decentralized, the losers of the vote could simply build a parallel space. The Cabal could not stop `alt.*`; it could only decline to carry it. This is self-governance working, and it is also self-governance showing its limit. The same property that let dissenters escape the Cabal would later let others escape every norm the community had.

7 Cooperation, Free-Riding, and the Limits of Self-Governance

Why did a system with no enforcement work at all? Kollock and Smith [11] give the sharpest answer by treating UseNet as a commons. What was shared was not pasture or fish but bandwidth and, more importantly, attention. A newsgroup was useful only so long as its members contributed signal (answering questions, staying on topic, reading before posting) and refrained from consuming more than their share. As in every commons, the incentives cut against cooperation. Answering a stranger’s question costs you time and benefits mostly others; posting an off-topic rant costs you nothing and imposes a small cost on thousands. The rational individual free-rides, and if everyone does, the commons collapses.

Drawing on Elinor Ostrom’s studies of real-world commons, Kollock and Smith ask which conditions for successful self-governance UseNet met. Several, it turns out. Groups were small enough that regulars recognized one another. Identities were persistent, so reputation mattered and a poster who abused a group paid in standing. Norms were explicit, in FAQs and in documents like Shea’s, and enforcement was cheap: public rebuke, ridicule, or the *killfile*, which let a reader silently filter out an offender. For a community of a few thousand mutually visible people, this was enough.

The same analysis shows where the cracks lay, and shows that they were structural rather than accidental. Ostrom’s first condition is clearly defined boundaries, and UseNet had none: anyone could join, and no one could be excluded. Sanctions did not scale: a killfile protects one reader, not the group, and ridicule only works on people who care about the group’s opinion. Above all, the whole arrangement depended on newcomers arriving slowly enough to absorb the norms before they outnumbered those who held them. Remove any of these conditions and cooperation, however heartfelt, becomes unsustainable. UseNet was about to lose all of them at once.

8 The Break: Green Card Spam and the Signal-to-Noise Collapse

On April 12, 1994, two immigration lawyers from Arizona, Laurence Canter and Martha Siegel, used a small script to post the same advertisement, offering help with the U.S. “Green Card Lottery,” to roughly six thousand newsgroups at once [3]. It did not matter that the groups were about physics, or knitting, or Estonian culture; every reader of every group saw it. It was the first mass commercial spam, and the reaction was instant and furious: tens of thousands of complaints and the lawyers’ Internet provider knocked offline by the flood. When Canter and Siegel ran the same script again that June, a Norwegian programmer, Arnt Gulbrandsen, “hacked together a canceller in about a minute” and had cancels flowing within the hour, by his own account posted two days later [7]: the first *cancelbot*, which issued forged cancel messages to erase the ad wherever it appeared.

The community had, in a sense, won. The ad was deleted and the couple widely reviled. But look closer. Canter and Siegel were unrepentant, claimed the episode had earned them clients, and wrote a book on how to do it again. The cancelbot was itself a hack, exploiting the same absence of authentication that had made the spam possible, and it opened an arms race others soon joined. Every rule in *Netiquette* had been broken, and there turned out to be no mechanism whatever for enforcing any of them against someone who did not care about the community's regard.

Why, then, did UseNet actually fail? Rule out the easy answers first. Competition from centralized alternatives played at most a minor role; the Web was still tiny in 1994. UseNet was not too slow or insufficiently interactive: mailing lists, which prospered, were no faster. Legal and regulatory pressure was not the cause. Nor was an inability to monetize, or to scale technically; UseNet carried volumes that would have been unthinkable a decade earlier.

In the instructor's experience and potentially-biased perspective, the decisive factor was spam, and specifically the spiral in signal-to-noise ratio it set off. The Green Card incident was a moment; the condition it revealed was permanent. Once it was clear that anyone could post anything to everyone at no cost, the trickle of commercial and hostile noise became a flood, and the effort needed to find the worthwhile posts in any group rose year by year. People rarely leave a place because of one bad day. They leave when the daily cost of staying exceeds the benefit, and that is what began to happen. Readers scurried for better-protected venues: first mailing lists, with a moderator and a membership roll, and eventually the many commercial platforms that arose to promise exactly that protection. The next two sections follow that migration.

9 Eternal September

Spam attacked the commons from one side. From the other came something no one intended as an attack. UseNet had always had a September: each fall a new cohort of university students arrived, blundered through the norms, were corrected, and by October had become regulars. The system absorbed them because they came in one small wave a year, into a community that greatly outnumbered them.

In September 1993, America Online, then the largest consumer online service in the United States, began piloting a UseNet gateway with a limited beta group of subscribers, then opened it to its full membership in March 1994. The new arrivals did not come in a wave; they came continuously, in numbers that soon rivaled the existing population, and without the university culture that had informally tutored earlier cohorts. Veterans gave it a bitter name, dated to the pilot that started it all: the September that never ended, "Eternal September."

Keep this distinct from the Green Card spam of the following April. They were different phenomena, one innocent and one malicious, but they compounded the same underlying problem. Both raised the noise in every group faster than the group's norms could absorb it; both eroded the conditions Kollock and Smith identified. In Dahl's terms, Eternal September is a clean case of *enlightened understanding* breaking down at scale. The FAQs and accumulated knowledge were still there, but the fraction of participants who had read them fell toward zero, and every conversation restarted from nothing.

You will experience a small echo of this yourself. The UseNet server used for this course's first assignment is named Eternal September, a wry acknowledgment by its operators that what they keep alive is the network as it was after that month, not before.




Figure 1: Books published with “UseNet” in the title, by year. Source: Amazon.com and Library of Congress catalogs as of 29 August 2010.

10 Decline, Diaspora, and Legacy

UseNet did not end so much as recede. Figure 1 offers one crude but telling measure: the number of books published each year with “UseNet” in the title peaks in 1995, the year after the Green Card spam and two years after Eternal September, and falls off steeply thereafter. Peak UseNet as a cultural object coincides almost exactly with the moment its commons began to fail.

The people did not vanish; they moved, in the order the previous sections predict. First to mailing lists: smaller, subscription-based, often moderated, offering exactly the boundary UseNet lacked. Then, through the late 1990s and 2000s, to web forums and commercial platforms, from Slashdot and Yahoo Groups to Reddit and Facebook, which took UseNet’s core idea, topical many-to-many discussion, and wrapped it in a central authority that could ban, filter, and delete. Each inherited UseNet’s forms (threads, groups, moderators, reputation scores) and traded away its decentralization for a usable signal-to-noise ratio.

Smith’s empirical work [17], mining years of UseNet traffic, shows what remained: a network still large in raw volume but increasingly dominated by a few recurring roles and by automated posting. The symbolic endings came slowly. AT&T dropped its UseNet service in 2009 [1]; in December 2023 Google, which had hosted the largest public archive, stopped carrying new posts [6]. The network still exists, and you will use it in the first homework. But its role as the internet’s public square passed to systems whose custodians hold precisely the power UseNet was built to do without.

11 What Actually Happened: Revisiting the Promise

Return now to UseNet’s two roles, as a platform for publishing in Section 2 and for self-organization in Section 3, and to the questions that Sections 4 and 5 left open. With hindsight, what actually

happened?

Much of the promise was kept. As a publishing platform, UseNet gave ordinary people direct, many-to-many, gatekeeper-free communication on a global scale, for the first time in history, and for roughly fifteen years it worked. It produced real communities, real expertise freely shared, and much of the culture and vocabulary the internet still uses. As a platform for self-organization, its self-governance was not a fantasy: it rested on identifiable conditions, and while those held, it held. Anyone tempted to conclude that decentralized, ungoverned communication is simply impossible has to explain those fifteen years.

But the failures were not bad luck either, and this is the harder lesson. The properties that made UseNet democratizing (open access, no central authority, no cost to speak) were the same properties that made it defenseless. Dahl's criteria help name the pattern. UseNet was strongest on *effective participation*, and it was participation, unlimited and unfiltered, that eventually drowned *enlightened understanding*. It struggled with *control of the agenda* because it had no legitimate way to make collective decisions, so contested questions were settled by whoever could route around whom. And its weakness on *inclusiveness* was, uncomfortably, part of why it worked: a small, homogeneous, mutually visible population is easy to govern by norms alone. When UseNet finally became inclusive, it lost the conditions for self-governance and had nothing to replace them.

Licklider and Taylor's predictions sort the same way. They were right about communities of common interest, which UseNet realized more fully than anything before it; right about a shared library of knowledge; and right, in OLIVER, about an agent that now exists. They were early, at best, on work: unemployment did not disappear, and their worry about "a privilege or a right" was answered, for UseNet's entire life, "privilege." And whether communicating through a machine proved *better* than face to face is exactly the question the signal-to-noise collapse answers: better while the commons held, worse once it did not. UseNet never got its OLIVER, and what it lacked was precisely a way to filter.

The structural lesson is not that the promise was false but that it was incomplete. A technology can genuinely expand who may speak while doing nothing about who is heard, how decisions get made, or how a community defends itself against those who do not share its goals. Those problems do not solve themselves, and when a community cannot solve them, someone else, usually a company, eventually will, on its own terms. That trade is the story of the platforms that followed.

Throughout this course we will repeatedly be asking the same questions of a different technology. Which democratic goods did it genuinely deliver? Which conditions did that delivery depend on? And what happened when those conditions changed?

12 Looking Ahead

Two threads from this chapter continue directly. The first is moderation. Gillespie [5] traces the content-moderation dilemmas of today's platforms straight back to UseNet's struggles with spam, cancelbots, and the limits of norms. When this course reaches social bots, Sybil attacks, and the politics of identity and anonymity, remember that those platforms are attempting, with central authority and at enormous scale, exactly what UseNet attempted without either. The second thread is Licklider and Taylor's OLIVER, the automated agent that would read, filter, and respond on a person's behalf. That aside, made half in jest in 1968, becomes the central question of the chapter on AI and democratic society: what does democratic communication mean when not all the participants are people?

Further Reading

The readings assigned alongside this chapter, Licklider and Taylor [12] and Dahl [4], supply its two analytical frames; Hauben and Hauben [8] remains the fullest primary-era history for those who want it. For discussion groups, six shorter readings each take a different lens on the same story: Kollock and Smith’s commons analysis [11]; Ronda Hauben’s personal retrospective [9]; Smith’s quantitative study [17]; the *Columbia Journalism Review*’s account of the Green Card spam [3]; Slashdot’s popular retrospective [16]; and Shea’s *Netiquette* [15]. Rheingold [14] and Kollock and Smith’s edited volume [10] offer deeper background.

References

- [1] AT&T *Jettisons the Last of Its Usenet*. The Register. 2009.
- [2] John Perry Barlow. *A Declaration of the Independence of Cyberspace*. 1996. URL: <https://www.eff.org/cyberspace-independence> (visited on 09/01/2026).
- [3] Columbia Journalism Review. *Original Sin*. 2004. URL: https://www.cjr.org/critical_eye/original_sin.php (visited on 09/01/2026).
- [4] Robert A. Dahl. “Criteria for a Democratic Process”. In: *Democracy and Its Critics*. Yale University Press, 1989.
- [5] Tarleton Gillespie. *Custodians of the Internet: Platforms, Content Moderation, and the Hidden Decisions That Shape Social Media*. Yale University Press, 2018.
- [6] *Google Ends Usenet Links*. The Register. Dec. 2023. URL: https://www.theregister.com/2023/12/18/google_ends_usenet_links/ (visited on 09/01/2026).
- [7] Arnt Gulbrandsen. *Re: Now comes the C&S crunch... let’s see*. Archived copy via the Internet Archive Wayback Machine; the original Google Groups page is no longer available. Usenet post to `news.admin.misc`, Message-ID <2to5ck\$kiu@trondviggo.nvg.unit.no>. June 16, 1994. URL: <https://web.archive.org/web/20120618062221/https://groups.google.com/group/news.admin.misc/msg/2e7ce912be3c373f?dmode=source&hl=en> (visited on 09/07/2026).
- [8] Michael Hauben and Ronda Hauben. *Netizens: On the History and Impact of Usenet and the Internet*. IEEE Computer Society Press, 1997.
- [9] Ronda Hauben. *Remembering Netizens*. 2022. arXiv: 2210.02978. URL: <https://arxiv.org/abs/2210.02978> (visited on 09/01/2026).
- [10] Peter Kollock and Marc A. Smith, eds. *Communities in Cyberspace*. Routledge, 1999.
- [11] Peter Kollock and Marc A. Smith. “Managing the Virtual Commons: Cooperation and Conflict in Computer Communities”. In: *Computer-Mediated Communication: Linguistic, Social, and Cross-Cultural Perspectives*. Ed. by Susan C. Herring. John Benjamins, 1996. URL: <https://users.ece.utexas.edu/~perry/education/382v-s08/papers/kollock.pdf> (visited on 09/01/2026).
- [12] J. C. R. Licklider and Robert W. Taylor. “The Computer as a Communication Device”. In: *Science and Technology* (1968).

- [13] LivingInternet. *Netnews History*. URL: https://www.livinginternet.com/u/ui_netnews.htm (visited on 09/01/2026).
- [14] Howard Rheingold. *The Virtual Community: Homesteading on the Electronic Frontier*. 2nd ed. 2000, MIT Press. Full text freely hosted by the author at <http://www.rheingold.com/vc/book/>. Addison-Wesley, 1993.
- [15] Virginia Shea. *Netiquette*. Excerpt widely mirrored as “The Core Rules of Netiquette.” 1994. URL: <https://www.mccc.edu/~virtcoll/Netiquette.pdf> (visited on 09/01/2026).
- [16] Slashdot. *The Rise and Fall of Usenet*. Dec. 2023. URL: <https://tech.slashdot.org/story/23/12/20/2147253/> (visited on 09/01/2026).
- [17] Marc A. Smith. “You Are Who You Talk To: Detecting Roles in Usenet Newsgroups”. In: *From Usenet to CoWebs: Interacting with Social Information Spaces*. Ed. by Christopher Lueg and Danyel Fisher. Springer, 2003.