

December 12, 2020

Internet Standard-Setting and Multistakeholder Governance

Nick Doty

UC Berkeley, School of Information

Status of This Document *This is a chapter of a published dissertation: [Enacting Privacy in Internet Standards](#).*

1 Internet Standard-Setting and Multistakeholder Governance

This work takes standard-setting as the site for exploration of how basic values (particularly privacy and security) are considered and developed in the design and implementation of large-scale technical systems.

Standards are the kind of unthrilling artifacts that are often taken for granted, assumed as a background quite separate from the concrete technologies themselves. When we think of the history of the railroad, for example, we are more likely to remember the rail magnates or the massive construction of the transcontinental railroad rather than the debates over gauges, even though compatibility of rail gauge has important implications for transit design to this day. Technical standards are a kind of infrastructure, both essential for development and often invisible to the casual observer.¹

Standards are dull in that they're:

- dry as reading material;
- unexciting (typically) in the day-to-day practice of their development; and,

¹h/t Richmond Wong

- only indirectly connected to the implementation of Web technologies.

Standard-setting is, nonetheless, essential in that it's:

- required (practically and politically) for the development of Internet and Web functionality;
- impactful and lasting in its impacts, which may remain in use for years or decades;
- distinctive of the Internet and the Web compared to many other technological developments; and,
- where agreement between a wide range of stakeholders is worked out.

As described in this chapter, Internet and Web standard-setting uses an uncommon but practically-minded consensus process for decision-making, which has implications for legitimacy and interoperability. Because of the typically open and public process and unique structure at the boundary between organizations, standard-setting bodies provide a venue that is rich for study and a process that is potentially innovative. Finally, these multistakeholder groups, including individuals from various backgrounds and a wide range of sectors, represent a distinctive governance model of interest to policymakers around the world for addressing complicated, cross-border issues of public policy, including privacy.

1.1 What is a standard

In discussing Internet and Web standards, I should explain what a standard actually is in this context.

1. Standards are, often long, documents.
2. Standards define what a piece of software needs to do in order to be compliant with the standard and in order to work with other software.
3. Standards don't define anything else.

Standards are documents, rather than code. Web and Internet standards are typically written in English, but they rely heavily on technical language, precise terminology referring to particular definitions, ordered lists of steps to define algorithms, and in some cases formal syntax (like ABNF ([Overell and Crocker 2008](#)) or WebIDL (["Web IDL" 2018](#))).

These documents explain how a piece of software that implements that particular standard needs to behave. So, for example, the HTML standard describes how a

The image shows a vertical line of text, which is a table of contents for HTML5. The text is extremely small and dense, making it completely illegible. It appears to be a list of topics or sections, but the individual items cannot be read.

Figure 1: The table of contents for HTML5 (Moon et al. 2017). No really, this is just the table of contents, none of the actual content.

Web browser should represent an HTML page and its elements, and describes how the author of a Web page should use HTML markup for a document. HTML is a complicated language, enabling a wide range of documents and applications, and interacting with many other separate standards that define presentation and other functionality. Printed out, the HTML specification would be about 1200 pages long, with the first 20 pages just a table of contents.² Most users of the HTML standard won't ever print it out or have any need to read it at length, but it is an invaluable reference for developers of browser software.

When standards are present (whether they're *de facto*, *de jure*, or otherwise broadly adopted), interoperability is possible. You can plug a phone line into a port in your wall and into your home phone, and expect it to fit and to work the same for making calls, even though the manufacturer of the phone didn't manufacture the cable you used or install the plug in your wall. When you visit your hometown newspaper's website, you can (hopefully) read the articles and see the photos whether you're using Firefox, Edge, Safari, Chrome, Opera or UC Browser, and your newspaper's web editor probably hasn't even tested all of those.

To be precise, specifications uses *normative* language to define exactly the requirements necessary to be a conformant page or a conformant user agent (for example, a Web browser on a phone or other computer). Language like MUST, SHOULD, MAY, REQUIRED and OPTIONAL have specific meaning in these standards (Bradner 1997). Non-normative sections provide context, explanation, examples or advice, but without adding any further requirements. Standards are specific about those requirements in order, perhaps counterintuitively, to enable diversity. For every functional difference not normatively specified, different implementations can do different things – pages can be constructed in different ways, browsers can render pages differently, within different user interfaces, different privacy settings, different performance characteristics, with various tools for their users. Interoperability of implementations allows for diversity and if variation were not a desired outcome, no standard would be necessary: a common implementation would be sufficient, and much more efficient to develop than setting a standard.

1.1.1 Standards terminology This text will occasionally use “specification” and “standard” almost interchangeably, which is common in this area. However, a

²There is no canonical print version, but, for example, WHATWG publishes a PDF that could be used for printing: <https://html.spec.whatwg.org/print.pdf>. See the figure for a screenshot of the W3C HTML Recommendation's table of contents.

specification (or, “spec”) is typically any document setting out how a piece of software should operate, whether or not it’s stable, implemented, reviewed, accepted as a standard or adopted. A standard is a specification that has either a formal imprimatur or actual demonstrated interoperability. People write specifications, and hope they become standards.

“Standard” itself is a heavily overloaded term; it is used in distinct if related ways in different fields and settings. For one confusing example, economists sometimes refer to a dominant market position as a standard, as in the 1990s when Microsoft’s Internet Explorer appeared likely to become the standard. In that case, the standard of having a dominant market position actually inhibited interoperability or the development of the interoperable specifications we call Web standards. And standards are often described as some bar of quality or morality: regulations might set out performance standards as requirements on a regulated group that can be met in different ways or profane or otherwise inappropriate content may be restricted by the Standards and Practices department of a broadcaster (Dessart n.d.).

1.2 The consensus standard-setting model

We reject: kings, presidents and voting. We believe in: rough consensus and running code. — Dave Clark, 1992

Technical standard-setting is a broad field, encompassing a wide range of technologies and organizational models. This research looks primarily at the consensus standard-setting model, which is the typical approach for design of the Internet and the Web. Consensus standard-setting is particular to situations of voluntary adoption, as opposed to *de jure* standards set in law or through some authoritative commitment (Cargill 1989). Voluntary standards are in contrast to regulatory standards: where governments intervene in setting mandatory requirements, often on safety or necessities for an informed consumer. Cargill appears skeptical of regulatory standards that are too broad in scope or too antagonistic to industry as being difficult to enforce, with OSHA the primary example (1989). But he lists different strengths and weaknesses for voluntary and regulatory standards: in short, that voluntary standards have flexibility and support of industry adopters, while regulatory standards can more easily be centralized and enforceability is more feasible.

The phrase “rough consensus and running code” should be considered in contrast to consensus as it might be defined in other political contexts. This isn’t

typically operated as unanimous agreement, as some might understand “coming to unity” in the Society of Friends, for example, or a super-majority vote as the modified consensus of Occupy Wall Street assemblies was often operationalized. Instead, guided by implementability and pragmatism, these standards groups look for a “sense of the room” – often evaluated through humming or polling rather than voting. Consensus decision-making can be slow and frustrating, but it may also create a process for sustainable resolution (Polletta 2004).

As a practical matter, voluntary standards need to be broadly acceptable in order to be broadly implemented. But that practical intent also has important implications for the procedural and substantive legitimacy of standard-setting. Froomkin (2003) argues that Internet standard-setting approaches a Habermasian ideal of decision-making through open, informed discussion. While consensus Internet standard-setting may boast procedural advantages uncommon to many governance processes (around transparency and access in particular, even though barriers continue to exist in both areas), evaluating the substantive legitimacy additionally requires looking at the outcome and the ongoing relationship among parties (Doty and Mulligan 2013).

1.2.1 History of standards Cargill traces a long history of standards, starting with examples of language and common currency, and focusing on the enabling effect that standardization has on trade and commerce (1989). Standard measurements and qualities of products make it easier to buy and sell products with a larger market at a distance, and standardized rail gauges made it possible to transport those goods. Industrialization is seen as a particular driver of voluntary standards to enable trade between suppliers: standardized rail ties make it possible to purchase from, and sell to, multiple parties with the same product (Cargill 1989). A similar motivation affected the development of Silicon Valley, where computer makers preferred to have multiple chip manufacturers as suppliers, and each with multiple customers, to build stability in the industry as a whole (Saxenian 1996).

Information technology standards have some important distinctions from the industrial standards that we identify as their predecessors. While concrete precision was a prerequisite for measurement standards or the particular shapes and sizes of screws or railroad ties, software involves many abstract concepts as well as technical minutiae. And information technology also expects a different rate of change compared to more concrete developments. The slowness of developing consensus standards for the Internet presents a challenge and encourages the use of more nimble techniques (Cargill 1989, among others).

In many ways, voluntary Internet standards make up a common good – usable by all. As an economic matter, Internet standards have important distinctions from rivalrous goods. Where Ostrom defines commons and ways of preventing overuse of a pooled resource (2015), Simcoe describes “anti-commons” and encouraging adoption of a common technical standard (2014).

Like many collective action problems, developing open technical standards may suffer from free-riding. As Ostrom (2015) puts it:

Whenever one person cannot be excluded from the benefits that others provide, each person is motivated not to contribute to the joint effort, but to free-ride on the efforts of others. If all participants choose to free-ride, the collective benefit will not be produced. The temptation to free-ride, however, may dominate the decision process, and thus all will end up where no one wanted to be. Alternatively, some may provide while others free-ride, leading to less than the optimal level of provision of the collective benefit.

If the standard will be made freely available, unencumbered by patents or even the cost of reproduction, and any vendor is encouraged to use it, there may be a disincentive to investing time, money and effort in participation to produce more standards, or update standards, since your competitors get all the same benefits without the costs. However, as Benkler points out, these information goods don't require collective action regarding allocation (since copying and distributing a standards document has minimal costs and the resource doesn't get “used up”) and the larger number of users might actually increase the benefits of participation (2002).

At the same time, technical standards provide network effects: if they're widely adopted they can become market standards, locking in technology that will subsequently be used by other market players and applications that depend on those standards. So participation can itself be motivated by rent-seeking behavior, and competition between standards. As Simcoe notes, standard-setting bodies have developed some organizational methods to respond to these concerns.³

1.2.2 The Internet and Requests for Comment I don't have the expertise to provide a history of the Internet, nor is another history of the Internet needed.

³For more discussion of the economics, organizational structure and legal implications of standard-setting, see “[Legal considerations in standard-setting](#)” below.

However, in understanding how the Internet standard-setting process functions, it is useful to see the motivations and context in which it began and how the Internet has evolved from an experimental project into a massive, complex piece of infrastructure.

Where should one read for an Internet history? A small, non-exhaustive list of suggestions:

- Abbate's *Inventing the Internet* (2000) is a very readable history, including a detailed accounting of the development of packet switching, and the motivations for its use.
- Mathew traces the history more briefly, but with a particular focus on the social contexts: institutions and social relationships (2014, "A Social History of the Internet").
- Several people instrumental in the early Internet architecture have also written their own brief history of the Internet (Leiner et al. 2009).

The Internet is a singular, global network of networks, characterized by routing of packets and (mostly) universal addressing. Devices (laptops, phones, large server farms) connected to the Internet can communicate with one another, despite running different software and being connected to different networks, and use a wide range of applications, including telephony, email, file transfer, Web browsing and many more.

Among the earliest clearly identifiable forerunners of the Internet we know today was ARPANET, a project of the Advanced Research Projects Agency (ARPA), which we now know as the Defense Advanced Research Projects Agency (DARPA). Motivated by the goal of more efficient use of the expensive computational resources that were used by different ARPA projects located at universities and research centers, the agency supported research into networking those large, rare computers. The technology of packet switching had been suggested independently by different researchers both for fault tolerance (including, as is often cited, the ability for command and control networks to continue to function after a nuclear strike) and for remote interactivity (allowing multiple users of a remote machine in interactive ways). Packet switching provided an alternative to dedicated circuits, a more traditional design making use of telephone lines.

Graduate students at a few research universities were tasked with defining protocols for these remote communications. Those informal meetings, notes and correspondence eventually became the Network Working Group (NWG). The tentative uncertainty of those students – now known as the original architects of

the Internet – is well-documented, as in this recounting from Steve Crocker, the first RFC editor (2009):

We thought maybe we'd put together a few temporary, informal memos on network protocols, the rules by which computers exchange information. I offered to organize our early notes.

What was supposed to be a simple chore turned out to be a nerve-racking project. Our intent was only to encourage others to chime in, but I worried we might sound as though we were making official decisions or asserting authority. In my mind, I was inciting the wrath of some prestigious professor at some phantom East Coast establishment. I was actually losing sleep over the whole thing, and when I finally tackled my first memo, which dealt with basic communication between two computers, it was in the wee hours of the morning. I had to work in a bathroom so as not to disturb the friends I was staying with, who were all asleep.

Still fearful of sounding presumptuous, I labeled the note a “Request for Comments.”

The early networking protocols documented in those informal Requests for Comments (RFCs) were later supplanted by design and adoption of the Transmission Control Protocol and Internet Protocol, commonly considered together as TCP/IP.⁴ Driven in part by interest in network connections different than phone circuits, including radio communications to connect Hawaiian islands and satellite connections between seismic monitors in Norway and the US (Abbate 2000), these network protocols could be agnostic to the form of connection. All devices connected using these protocols, no matter what their physical connection or local network might be, could have individual IP addresses and reliable transmission of data (split up into packets and recombined) between them. This allows “internet-working”: communication between devices connected to different networks that are themselves connected.

While the the networking and internetworking protocols developed, the uses for ARPANET also changed. Originally designed for the sharing of access to large mainframe computers, many users preferred the communications capabilities.

⁴The design of these protocols is attributed to Vint Cerf and Robert Kahn, with the input and participation of many other stakeholders. TCP/IP is described in (1974) and RFCs 791 (1981a) and 793 (1981b).

Scientists shared data, programmers shared source code, and email unexpectedly became the most popular application on the ARPANET, including emails to the program managers who provided military and academic funding and early mailing list software for group discussion of topics of interest, like science fiction (Abbate 2000). Email, driven by the users, became an influence for developing shared networks for communications. And in using the tool of email to debate and construct an alternative architecture for the Internet, that community of users fits the concept of a “recursive public” (Kelty 2008).⁵

Organizationally, the Network Working Group gave way to the Internet Configuration Control Board, later replaced by the Internet Advisory Board, subsequently renamed the Internet Activities Board, which became popular enough to be subdivided into a number of task forces, most significantly the Internet Engineering Task Force and the Internet Research Task Force. The IAB changed names and tasks again to be the Internet Architecture Board, which still exists today, providing some expert advice and leadership to IETF tasks.⁶

While I have focused here on the development of Internet standards and the Internet standards process, this development did not happen in a vacuum. In parallel, computer manufacturers developed proprietary standards for networking their own devices. Telecommunications carriers, hoping to limit the power of these proprietary standards, developed network protocols that relied on “virtual circuits” where the network provided reliable communications. While packet switching expected “dropped” packets and different routing mechanisms and required hosts to handle those variations, the approach of circuits put the responsibility for reliable delivery on the network.

The International Organization for Standardization (ISO), a formal international standards organization operating with the votes of different representatives of standards organizations from each nation state, started the development of OSI network standards, in cooperation with the International Telecommunications Union Standardization Sector (ITU-T), an agency of the United Nations that had been developed to set cross-border telegraph and telephone standards. The OSI work included the still influential seven-layer networking model, as well as standards to implement those different layers. Like many questions of standards adoption, various economic and political factors come into play: the

⁵Kelty specifically concludes that the Internet itself is not a recursive public, but the technical contention over the ARPANET, NSFNET and early Internet may be a closer fit for the concept. See, later, “Ethnography in Internet Standard-Setting” for more discussion of this concept.

⁶“A Brief History of the Internet Advisory / Activities / Architecture Board” (n.d.) documents the history of these confusing name and abbreviation changes.

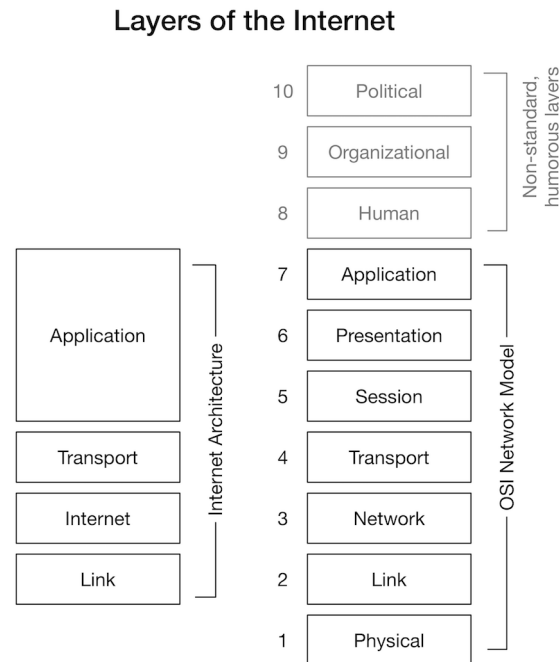


Figure 2: Layers of the Internet, both the OSI seven-layer model and the TCP/IP four-layer model (Braden 1989), aligned.

relatively wide deployment and military use of TCP/IP in ARPANET, European government support of ISO standards to provide a common market for technology across European countries, the relative market powers of computer manufacturers, telecommunications carriers and Federally-funded universities and research centers, the timing of releases of competing standards (Maathuis and Smit 2003; DeNardis 2009).

From an IETF participant's perspective, ISO's process was long and complicated, and the standardized protocols were lacking in widespread implementations. While OSI protocols might have had some potential advantages (in areas of security, or the size of address space), that TCP/IP was running and working, freely available and already implemented, were more germane. Being simple and just good enough to work would become common advantages of the relatively informal IETF model. When the IAB, a smaller group of technical leaders, made a proposal to adopt the OSI CLNP protocol as the next version of the Internet Protocol, there was widespread anger from IETF participants at the possibility of

top-down development of protocols or switching to the more formal ISO process. It was in response to this concern that Dave Clark made his famous description of IETF's "rough consensus and running code" maxim.

IETF's process today is a little more formal than its origins, but retains many informal characteristics. Leadership on technical standards is provided primarily by the Internet Engineering Steering Group (IESG) a rotating cast of volunteer Area Directors (ADs), selected by the Nominating Committee (NomCom), which is itself drawn from regular meeting attendees. The Area Directors make decisions on chartering new Working Groups, a process involving an informal "birds of a feather" meeting to gauge community interest, recruiting chairs to manage the work and gathering feedback on a charter of the group, its scope and deliverables.

IETF Working Groups can be operated in different ways, but often follow a similar model. The appointed chairs have significant authority to manage the group's work: setting the agendas for meetings and foreclosing topics out of scope, selecting editors to develop specifications, and determining the consensus of the group for decision-making purposes. Discussion happens most often on publicly-archived mailing lists, with in-person meetings as part of the three-times-a-year IETF meeting schedule (and for some very active groups, interim in-person meetings between the IETF meetings). While in-person meetings can be significant venues for hashing out issues, all decisions are still confirmed on mailing lists.

The IETF does not have any formal membership, for individuals, organizations or governments. This lack of membership has some distinctive properties: for example, it makes voting largely infeasible. Participation is open to all, by engaging on IETF mailing lists or attending in-person IETF meetings.⁷ The lack of organizational membership also contributes to the convention that individuals at IETF do not represent or speak for their employers or other constituents; instead, individuals speak only for themselves, typically indicating their affiliations for the purpose of transparency.⁸

Attendees at particular IETF meetings pay to defray some meeting costs and companies pay to sponsor those meetings, but remote meeting participation and participation on mailing lists does not incur any fee. The activities necessary to operate the IETF are largely supported by the employers of its volunteers, but paid

⁷As a result, just counting the number of participants in IETF's work is challenging. We are exploring some such measures via automated mailing list analysis: see [this notebook on IETF participation](#) and [this presentation on IETF mailing list analysis](#).

⁸The tradition of individual participation is considered in more detail in [Individuals vs organizations in standard-setting process](#).

staff and other costs are funded by the Internet Society, whose major budget now comes from the sale of .org domain names.⁹

The RFC series began with that note from Steve Crocker on the protocols for ARPANET host software; each is numbered, with that first one considered RFC 1. Today, RFCs are more vetted than a simple request for comments, but come from different streams and have different statuses, representing maturity or purpose. The review of the IESG is necessary for publishing a document as an RFC, with different requirements for different document types, but typically requiring the resolution of any significant objections. Such objections are called a `DISCUSS` and, fitting the name, are designed to promote finding an alternative that addresses the objection, rather than a direct refusal.

Of over 8000 RFCs, only 92 have reached the final level of Internet Standard. For example, STD 90, also known as RFC 8259, describes JSON, the JavaScript Object Notation data format, in widespread use. Over 2400 are “informational” and 400 more are “experimental”: these are RFCs that are not standards and aren’t necessarily intended to be, but document some technique for consideration, some protocol that may be used by some vendors, or some documentation of problems or requirements for the information of readers. These vary significantly, but, for example, RFC 6462 reports the results of a workshop on Internet privacy; RFC 1536 described common problems in operating DNS servers. Other RFCs are not Internet technology specifications at all, but guidance on writing RFCs or documentation of IETF meeting practices: RFC 3552 provides advice to document authors regarding security considerations; RFC 7154 describes a code of conduct for participation in IETF; RFC 8179 sets out policies for patent disclosures.

That an RFC can be a request for comments, a well-established Internet standard, an organizational policy or a particular vendor’s documentation, all with sequential numbers, can be confusing. RFC 1796 “Not All RFCs are Standards” was published in 1995 noting that topic, and the discussion continues with “rfc-plusplus” conversations. But RFCs remain diverse: they can be humble, informational, humorous, experimental; they are all freely available and stably published in good old-fashioned plain text; and, sometimes, they are established Internet Standards.

1.2.3 The Web, Recommendations and Living Standards Though commonly confused, the Web is distinct from the Internet; it is an application built on top

⁹Funding was less steady prior to ICANN’s 2003 allocation of .org domains to the ISOC-created Public Interest Registry. ISOC had relied largely on company members to provide sponsorships and pay membership dues.

of the Internet. The Internet is that global network of networks that lets computers communicate with one another enabling all sorts of applications; the Web is a particular application that lets you browse sites and meaningful pages and applications at particular locations.¹⁰

Where should one read for a history of the Web?

- Robert Cailliau co-authored a book on the topic, *How the Web was born* (Gillies and Cailliau 2000)
- Tim Berners-Lee gave a “How It All Started” presentation, with pictures and screenshots, at a W3C anniversary (2004)

The World Wide Web began as a “hypermedia” project for information-sharing at CERN, a European research organization that operates particle accelerators in Switzerland. Developed by Sir Tim Berners-Lee and Robert Cailliau, among others, a protocol (HTTP), markup language (HTML) and client (the WorldWideWeb browser) and server (httpd) software made for basic functionality: formatting of pages and hyperlinks between them. This functionality was simple in comparison to hypertext proposals of the time, but the simple authoring and sharing of text and other resources combined with the connectivity of the Internet became an extremely popular application.¹¹

Web standardization was driven by the babel-style confusion of the “browser wars.” Inconsistencies meant that a page written using some features might look entirely different in one browser compared to another. Sites might include a disclaimer (and in some ways, a marketing statement) of, for example, “best viewed in Netscape Navigator 4.” This situation is a frustration for the reader and a challenge for the author. And affecting a wider range of market players (site authors, browser vendors, even Internet providers), it potentially undermines the use of the Web altogether.

The World Wide Web Consortium (W3C) was formed in 1994, hosted at the Massachusetts Institute of Technology, with Sir Tim Berners-Lee, the inventor most directly responsible for the Web and the Hypertext Markup Language

¹⁰This chapter won’t provide a detailed technical description of the Internet and the Web. Instead, see the system overview sections of [Encrypting the Web, a “handoff”](#) and [Do Not Track, a “handoff”](#).

¹¹“How many web pages are there?” is a simple, interesting and unanswerable question that’s asked from time to time. An imperfect measure: Google announced they had indexed a trillion pages in 2008, up from 26 million in 1998 (Alpert and Hajaj 2008).

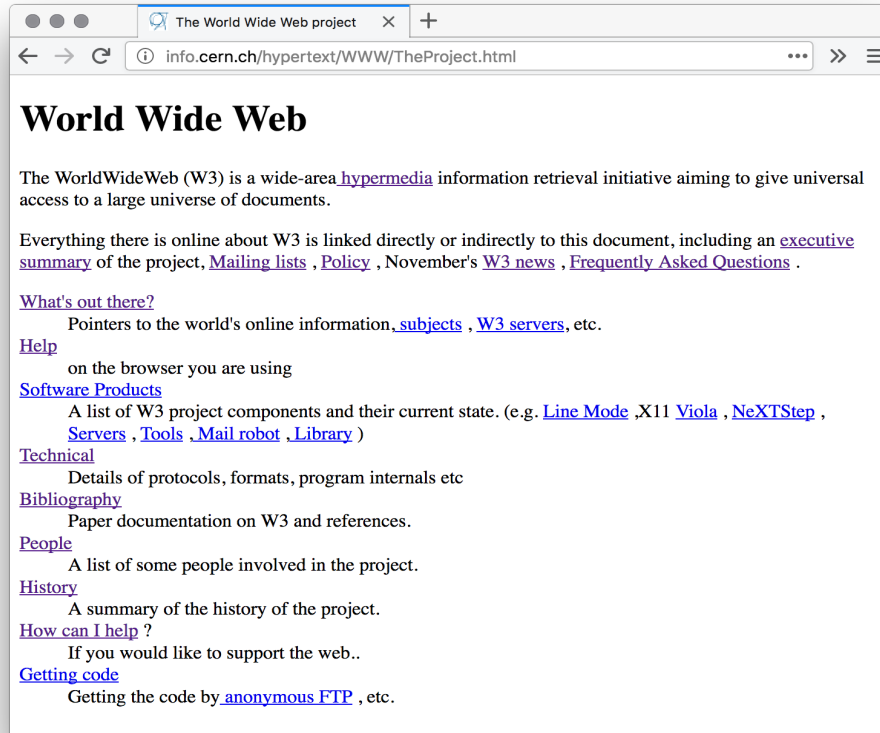


Figure 3: The first ever Web site is again operational on CERN's servers, with early descriptions of the Web, its operation and motivations.

(HTML), as its Director. HTML had a home, and, soon after, a process¹² for further development.

W3C's "consortium" model relies primarily on membership for funding¹³ and direction. Its 479 member organizations¹⁴ are mostly companies, with some universities, non-profit organizations and government agencies. Those companies are a mix of small, medium and large; they reach across industry sectors with,

¹²Or rather, a Process: <https://www.w3.org/Process>.

¹³Funding temporarily included support from the Internet Society (Jacobs 2009).

¹⁴That membership changes over time. 479 members as of 21 August 2018: <https://www.w3.org/Consortium/Member/List>.

as you might expect, a particular representation of technology-focused firms.¹⁵ W3C employs a staff (sometimes called “Team”) who coordinate work and handle administrative tasks, but the actual process of standardization is done by volunteers, most often those employed by member organizations, and the general direction of what work to do is set by the member organizations, who send representatives to an Advisory Committee.

Standards are developed by Working Groups: smaller groups (typically with 10 to 100 members), with a charter to address particular topics in specific deliverables. As of August 2018, W3C had 36 Working Groups actively chartered to address topics ranging from accessibility guidelines to the Extensible Stylesheet Language (XSLT).¹⁶ The documents that become standards follow an iterative process of increasing breadth of review and implementation experience: an Editor’s Draft is simply a document in progress, a Working Draft is published by a Working Group for review, a Candidate Recommendation is a widely-reviewed document ready for more implementation experience, a Proposed Recommendation has demonstrated satisfaction of all requirements with sufficient implementation experience and a Recommendation shows the endorsement of W3C membership (fantasai and Rivoal 2020).¹⁷ That the most complete and accepted stage of a technical report is a “Recommendation” emphasizes the humility of this voluntary standards process (not unlike “Request for Comment”) – even a published Recommendation doesn’t have to be adopted or complied with by anyone, even W3C’s members, even the members of the Working Group that worked on it, even the employer of the editor of the document. It’s just that, a recommendation.

Working Groups at W3C can operate using different procedures but typically follow a similar process, guided by the collective advice of past participants (“The Art of Consensus: A Guidebook for W3c Group Chairs, Team Contact and Participants” n.d.). An editor or group of editors is in charge of a specification, but key decisions are made by consensus, through discussion by the group in meetings, teleconferences, email and other online conversations and as assessed by

¹⁵The overlapping stakeholder groups at W3C figure in the [Methods chapter](#) maps out a rough sense of the stakeholder groups and member groups represented in W3C. Quantitative analysis of the member organizations is possible, but not included here – crowdsourcing proved challenging and the process is tedious. However, some work on this is underway as part of the ongoing study of civil society organization participation in Internet governance by the University of Exeter: <http://www.internetpolicystreams.com/>.

¹⁶W3C maintains a list of current and past groups: <https://www.w3.org/Consortium/activities>.

¹⁷The exact details of these stages of review have changed over time, but the iterative process of increasing review and experience has remained consistent.

the chairs who organize the group's activity.¹⁸ This process aims for sustained objections to a group's decisions to be uncommon, but processes for appealing decisions are in place. The Director plays an important guiding role in addressing objections and evaluating maturity, but decisions can also be appealed to a vote of the membership.

As new standardized versions of HTML were published at W3C, a split grew between XHTML – a set of standards that some thought would enable the Semantic Web and XML-based tooling among other things – and updating versions of HTML that instead reflected the various document and app uses of the Web. The Web Hypertext Application Technology Working Group (WHATWG)¹⁹ formed in 2004 from browser vendors (specifically, Apple, Mozilla and Opera) who wanted to update HTML with application features that were under development rather than pursuing an XML-based approach. Work on subsequent versions of XHTML was dropped and W3C and WHATWG processes worked in parallel on HTML5, published as a W3C Recommendation in 2014. Tensions remain between W3C and WHATWG and supporters/antagonists of each, but the work of technical standard-setting continues in both venues – on HTML, which is published both by WHATWG as a Living Standard and as a versioned document at W3C,²⁰ and on other specifications. Paul Ford's description in *The New Yorker* is accessible, and, to my eyes, remains an accurate assessment (2014):

Tremendous flareups occur, then settle, then threaten to flare up again.
[...] For now, these two organizations have an uneasy accord.

WHATWG has a distinct process for developing standards, although there are many similarities to both IETF and W3C process, and those process similarities have increased substantially with a new governance and IPR policy agreed upon in late 2017 (van Kesteren 2017), with the formal inclusion of Microsoft in the process.

¹⁸The day-to-day details of this process are discussed further in [A Mixed-Methods Study of Internet Standard-Setting](#).

¹⁹Why the strange, long acronym? Apocryphally, because it started as this secretive separate process and it seemed like a good joke to be able to say, in response to a question like, “are you working with some other rival working group?” “what working group?”

²⁰As of May 2019, WHATWG and W3C have explicitly agreed on a Memorandum of Understanding with the goal of a unified HTML specification, still with both Living Standard and versioned, reviewed snapshots (“[Memorandum of Understanding Between W3c and WHATWG](#)” 2019).

Discussion in WHATWG happens primarily on GitHub issue threads and IRC channels (and, in the past, mailing lists) and in-person meetings are discouraged (or, at least, not organized as WHATWG meetings) for the stated purpose of increasing the breadth of access (“FAQ — WHATWG” n.d.). While W3C and IETF use versioned, iteratively reviewed documents with different levels of stability, WHATWG publishes Living Standards, which can be changed at any time to reflect new or revised features. (However, as of late 2017, fixed snapshots are published on a regular basis to enable IPR reviews and patent exclusion, similar to the W3C process.) Rough consensus remains a guiding motivation, but WHATWG implements consensus-finding differently, relying on the assessment of the Editor of each specification. The Editor makes all changes to each specification at their own direction, without any process for chairs or separate leadership to assess consensus. (However, an appeals process for sustained disagreement is now in place, with decisions put to a two-thirds vote of the four companies that make up the Steering Group.) Because there is no formal membership (more like IETF’s model), there are not separate Working Groups, although there are Workstreams, which must be approved by the Steering Group, and all contributors must agree to a contribution agreement, which includes similar IPR commitments as in W3C Working Groups.

This research project primarily focuses on W3C and IETF standard-setting processes, although WHATWG and other groups may also be relevant at times. Other standard-setting bodies (or similar groups) also produce standards relevant to the Web and to privacy, often with either a narrower or broader scope. For example, the FIDO Alliance²¹ develops specifications for alternatives to passwords for online authentication; the Kantara Initiative²² publishes reports regarding “digital identity”; the Organization for the Advancement of Structured Information Standards (OASIS)²³ has a consortium model for standards on a wide range of information topics, particularly XML document formats and business processes, but have also worked on standards for privacy management and privacy-by-design. Broader still, the US government’s National Institute of Standards and Technology (NIST)²⁴ has a scope including all of science and technology, including specific process standards on privacy risk management (Brooks et al. 2017) and the basic weights and measures (among other things, they keep the national prototype kilo-

²¹<https://fidoalliance.org/>

²²<https://kantarainitiative.org/>

²³<https://www.oasis-open.org>

²⁴<https://www.nist.gov/>

gram), and the International Organization for Standardization (ISO)²⁵ welcomes national standard-setting organizations like NIST as its members, and covers an enormous scope from management standards for information security (“ISO/IEC 27001 Information Security Management” 2013) to “a method of determining the mesh-breaking force of netting for fishing” (“ISO 1806:2002 - Fishing Nets -- Determination of Mesh Breaking Force of Netting” 2002).

The divisions between W3C and WHATWG are useful to explore as a comparison regarding organizational policy: forum shopping is easier to see in such a direct side-by-side situation; that anti-trust, IPR and governance policies are apparently necessary for growing participation, especially for a large firm with an antitrust history as in the case of Microsoft, is more easily demonstrable. But the W3C and WHATWG models also invite comparison of different approaches to the Web and its standards.

Interoperable implementations are key to all the Internet standards processes discussed here, but WHATWG is especially specific about major browser implementations as the essential criterion guiding all other decisions. The model of a Living Standard reflects the increasingly short release cycles of different versions of those major browsers. For years, the “informed editor” distinction was especially contentious: Ian Hickson (known as Hixie) edited HTML in both the WHATWG and W3C processes, and decried certain decisions by the W3C Working Group contrary to his own as “political.”²⁶ While in many ways the informed editor approach is similar to the motivations behind other consensus standards body decision-making practices (decisions are not supposed to be votes, arguments are to be evaluated on their merits and implications, not on their loudness or how widely shared they might be), the apparatus of chairs, membership and governance/appeals processes add an element of represented stakeholders to decision-making, outside a singular technocratic evaluation.²⁷

Whether Recommendations or Living Standards, the Web’s protocols are defined in these Web-hosted documents and reflected in the voluntary, sometimes incomplete, mostly interoperable implementations in browsers, sites and other software.

²⁵<https://www.iso.org>

²⁶See this email thread from 2010 for example: <https://lists.w3.org/Archives/Public/public-html/2010Jun/0217.html>

²⁷These dual goals/modes will be an ongoing tension and opportunity. See, for example, [Individuals vs organizations in standard-setting process.](#)

1.2.4 Legitimacy and interoperability In evaluating the legitimacy of any decision-making process, including these rough consensus standard-setting processes, it may be useful to distinguish between procedural and substantive legitimacy. In the context of technical standard-setting, these have also been described as input and output legitimacy (Werle and Iversen 2006). In short, (1) are the steps of a process fair? and (2) is the outcome of the process fair to those affected?

Procedurally, we might consider access to participate meaningfully and transparency of decisions and other actions as hallmarks of legitimacy. The tools and practices common in Internet standard-setting can provide remarkable inclusion and transparency, while, simultaneously, substantial barriers to meaningful participation persist. On the one hand, anyone with an Internet connection and an email address can provide comments and proposals, engage in meaningful debate and receive a significant response from a standard-setting group. Anyone interested in those conversations at the time or after the fact can read every email sent on the topic, along with detailed minutes of every in-person discussion. On the other hand, discussions can be detailed, technical, obtuse and time-consuming, limiting meaningful participation to those with both the technical ability and the resources (time, money) to sustain involvement.

While we would anticipate that procedurally legitimate process is likely to be substantively legitimate as well, that might not be guaranteed: a majoritarian voting structure could seem legitimate while putting an unfair ultimate burden on some minority group, for example.

In consensus standard-setting, interoperability and voluntary adoption are the distinctive characteristics of success. Voluntary adoption may promote substantive legitimacy in some important ways: implementers and other adopters are not compelled to adopt something that they find out of the reasonable range, as we can see from the many completed technical standards that do not see widespread adoption. Engagement from stakeholders in design of a technical standard may encourage design of a workable solution for those stakeholders, rather than having a separate party (like a regulator or arbitrator) hand down a decision. But the success criteria of interoperable, voluntary adoption do not ensure the satisfaction of values-based metrics. In particular, stakeholders who are not themselves potential implementers – including government agencies or typical end users, say – have more limited opportunities to affect adoption, which might limit their influence on the substantive outcome. While interoperability may provide functionality and portability, that functionality may not meet users' needs or protect them from potential harms.

How procedural and substantive legitimacy may apply to the decisions of

consensus technical standard-setting processes, especially in technical standards with public policy importance, is detailed further in earlier work.²⁸ These same criteria will be especially relevant in comparing how the coordinating and decision-making function of standard-setting compares to other governance models (see [Drawing comparisons](#) below).

1.3 Organizational structure

1.3.1 How Internet standards bodies are structured As a matter of legal incorporation, Internet and Web standard-setting bodies have unusual structures. W3C is not a legal entity. WHATWG is not a legal entity. IETF is not a legal entity although, just recently,²⁹ there has been the creation of an LLC to provide a legal home for its administration. Until recently, none have had bank accounts of their own that can deposit checks, though IETF now will. Instead, W3C is a set of contracts between four host universities and the various member organizations; IETF is an activity supported by the Internet Society, a non-profit, and administered by a disregarded entity of the Internet Society; WHATWG is an agreement signed by four browser vendor companies.

Those legal minutiae are perhaps not the most germane consideration for the participants or for an analysis with organizational theory, but this structure (or lack thereof) is distinctive. Rather than independent entities, standard-setting bodies functionally exist through the activities of participants. Making that abstract concept real through analogy can be tricky, but, for example, one can think of the standard-setting body as a restaurant with tables around which people eat and talk ([Bruant 2013](#)). ISO describes itself as the “conductor” to an “orchestra [...] of independent technical experts” (“[We’re ISO: We Develop and Publish International Standards](#)” n.d.).

This may be an example of *institutional synecdoche*,³⁰ where there is confusion in distinguishing between the actions of an organization and of its component participants. When people complain about W3C (and people *love* to complain about W3C), are they typically attributing their complaint to W3C staff, or the documented W3C process, or the typical participants? There is certainly confu-

²⁸See [Doty and Mulligan \(2013\)](#), citing in particular [Tyler and Markell \(2010\)](#) on criteria for the acceptability of processes and [Lind and Tyler \(1988\)](#) for the social psychology of how participants perceive a process as procedurally legitimate.

²⁹As of August 27, 2018 ([Haberman, Hall, and Livingood 2020](#)), in the middle of drafting this chapter.

³⁰h/t Daniel Griffin, for the lovely term

sion about what these standards organizations are or what authority they have. For example, during a Senate committee hearing on the status of Do Not Track negotiations, there seemed to be genuine confusion among Senators over what W3C or its authority was, and why couldn't the different parties just find a room for discussions and coming to agreement, before it was pointed out that it was a voluntary process where companies were trying to come to agreement (Rockefeller 2013).³¹

There are other unusual organizational designs in Internet governance more broadly; for example, the IANA function has been a single person, a California non-profit under contract with the US Department of Commerce, and, post-transition, a non-profit absent government control. See [What is Internet Governance](#) below.

1.3.2 Standards are a boundary It can be tempting to conceive of the Internet and the Web as organizational fields, with the standard-setting bodies as sites where the field communicates, but the diversity of stakeholders and the diversity-enabling function of technical standards instead suggests understanding standard-setting bodies as boundary organizations.³²

Organizational fields can be defined in distinct ways, but consider DiMaggio and Powell's definition as a popular one: "those organizations that, in the aggregate, constitute a recognized area of institutional life: key suppliers, resource and product consumers, regulatory agencies, and other organizations that produce similar services or products" (1983). This includes elements of, but is not limited to, organizations that interact (connectedness) and companies that compete. Multistakeholder standard-setting does include some of these characteristics: organizations connect and communicate regularly through the standard-setting process, some of them are either competitors or have consumer/supplier relationships, and developing the Internet or the World Wide Web might be seen as a "common enterprise" (P. DiMaggio 1982).

³¹"Senator MCCASKILL: But I am a little uncomfortable that all of us seem to have agreed in the room that we are ceding the authority to set this policy to some organization I am not even sure who is in charge of this organization. Who do they answer to? Who are they, and how did we get to this point?" [...] "So what you are basically saying is this is just a place to go to try to see if all of you guys can agree? Couldn't we just set a room somewhere and all of you get there and try to decide and see if you all agree?" [...] And later, to laughter throughout the room: "Senator THUNE: Mr. Chairman, I would say that on behalf of a number of colleagues on my side that we would be really worried if W3C is run by the U.N."

³²This argument has previously been made in [Doty and Mulligan \(2013\)](#), but it is expanded here.

In other ways, though, participants in Web and Internet standardization demonstrate substantial diversity less characteristic of an organizational field. The Web browser vendors are certainly competitors, but their business models and corporate structures are quite distinct: Microsoft earns money largely through software sales, Apple through hardware sales, Google through online advertising, Mozilla is a non-profit, with revenue from search engine partners and donations. Most W3C members don't develop browsers: there are academics, consumer advocacy non-profits, Web publishers, retailers, telecommunications companies, online advertising firms and government agencies. Discussions can be tense when individuals from organizations in different industries interact and conflict: for example, online advertising firms, consumer advocates and browser vendors in the Do Not Track process or middlebox providers, financial services firms and client software developers in TLS. That standard-setting can be a difficult interpersonal process is known, but this work will explore some of those heightened tensions around privacy and security contestation.³³

In addition to the characteristics of the participants, the outputs of technical standard-setting bodies – that is, the technical standards themselves, give us some insight into the organizational structure because of their uncommon purpose. Technical standards, as described above, allow for flexibility by being specific about certain features of technical interoperability. They may qualify as “boundary objects” in the way that some STS scholars have described them: by providing interpretive flexibility of a single artifact (whether concrete or abstract), a boundary object allows for collaboration across different social worlds (Star and Griesemer 1989).

Rather than the site of an organizational field, we have identified these multi-stakeholder standard-setting bodies as boundary organizations (Doty and Mulligan 2013). The concept of “boundary organizations” was described by Guston in the specific context of the relationship between science and science policy. In order to both maintain the boundary between science and politics, but also blur that boundary enough to make connections across it to facilitate scientific-driven policy, Guston argues that boundary organizations can “succeed in pleasing two sets of principals” (2001). Three criteria define these organizations:

1. they enable the creation of boundary objects (or, related, “standardized packages”) that can be used in different ways by actors on either side of the boundary;

³³See [How standard-setting accommodates, succeeds and fails](#) in the findings.

2. they include the participation of actors on both sides, as well as a professional staff to mediate and negotiate the boundary;
3. they are accountable to both sides, politics and science.

The Office of Technology Assessment is a prominent and perhaps reasonably well-known example. While other advisory organizations were often considered partisan or co-opted, many saw the OTA as a respected and neutral source of analysis into technology and the impacts of policy proposals.³⁴ Its reports were boundary objects, in that they could be used by different committees or political parties for different purposes.

This early description of boundary organizations assumes exactly two sides: science and policy, or almost analogously, two political parties: Democrat and Republican. That bilateral, oppositional view seems to come from the particular literature of science and technology studies and Latour's view of science as Janus, the two-faced Roman god who looks into both the past and the future. The Janus metaphor is used in multiple ways, but most distinctively, it notes that science can simultaneously be seen as uncertainty – the practice of science involves a messy process about things that are by their nature not yet understood – and certainty – that science is what has already been settled and can be assumed (like a black box) for future work (Latour 1987).

But while it's tempting to see boundaries and conflicts as always two-sided, the concept of boundaries and boundary organizations can be applied more broadly. A particularly relevant description of boundary organizations comes from O'Mahony and Bechky, who describe how social movements that might be seen in direct conflict with commercial interests sometimes find success in re-framing objectives and maintaining collaborations where interests overlap. Boundary organizations allow for collaboration between organizations with very different interests, motivations and practices. In the case of open source software development, several open source projects have developed associated foundations to serve that boundary role: those foundations let corporations collaborate on the open source project by having a formal point of contact for signing contracts and representing project positions, without violating the openness practices of open source projects or requiring private companies to discuss all their plans in public (O'Mahony and Bechky 2008). Many of the other boundary management practices identified

³⁴Cf. attacks on W3C as "once neutral." Or political party attacks on the CBO when it scores their tax plans. The very shape of the attacks tells us something about the perceived position of each target organization.

related to individual rather than organizational control; open source contributors had reputation and impact on a particular open source project that followed them even when changing employers (O'Mahony and Bechky 2008). A similar ethos is present in Internet standard-setting, particularly, but not exclusively, at the IETF.³⁵

Internet standard-setting matches this definition of a boundary organization, but operates at an intersection of more than two clearly separable sides. Standards are boundary objects – agreed upon by different parties with some interpretive flexibility that can subsequently be used by different parties, including competitors and different sides of a communication. The multistakeholder standard-setting process involves participants from those diverse parties, with some professionals to help coordinate and mediate. And, ideally, these bodies are accountable to those different parties, whether that's users, different groups of implementers or even policymakers.

Even as we see WHATWG start to adopt much of the organizational structure of other Internet standard-setting bodies – a governance system, IPR rules, scoped working groups, etc. – it remains structured more like a field and less like a boundary. The steering group is limited to Web browser vendors (market competitors engaged in a collaborative common enterprise) and the guiding interoperability principle is browser vendor adoption, there is less indication of accountability to multiple, diverse principals.

A hypothesis to be explored or tested at a later date: if the WHATWG approach is a field rather than a boundary, then moving more standards to a WHATWG model should promote stronger forces of isomorphism among browser vendors. We could see the profession become “Web browser developers” rather than just “Web developers.”

This isn't an all or nothing situation – standards can also clearly be tools to enable supplier/consumer relationships and Web publishers and Web browser vendors can reasonably be seen in that light. The connectedness of a standards group can enable some of the professionalization and cross-pollination while also maintaining the distance of commercialism and non-profit/open-source activity.

How we classify standard-setting bodies (boundary vs. field) is not some academic exercise or merely a question of naming. Identifying the appropriate structure from organizational theory can let us apply insights from, and contribute learning back to, research into the sociology of organizations. In that very well-cited paper from DiMaggio and Powell, we see that fields typically exhibit forces (coercive, mimetic and normative) towards institutional isomorphism (1983) –

³⁵See *Individuals vs organizations in standard-setting process*.

we expect similar structures across the organizations, both as innovations are spread and as further diversification is restricted. Boundary organizations, in contrast, specifically enable collaboration among a diverse group and boundary objects can provide an interface for cooperation between groups that often have friction. Specifically, boundary organizations have been suggested as a kind of organizational method to allow social movements to collaborate with corporations and effect change.

As Colin Bennett describes (2010), privacy advocates have emerged in response to increasing surveillance, engaged in “collective forms of social action” and reflected in more common public protest to technological intrusion. While Bennett distinguishes this broad, networked activity from a worldwide social movement, there are certainly similarities in the diverse strategies and loose coalitions between numerous organizations and the dedicated individuals who participate. Privacy advocates practice in spaces beyond traditional non-profit advocacy organizations and also seek to work with or influence the behaviors of government and corporate actors.

Based on this model, the empirical work of this dissertation seeks to shed light on the following questions raised by this background. If Internet standard-setting organizations play the role of boundary organizations in mediating technical policy conflicts when it comes to Internet privacy and security, can they provide a way for privacy advocates to collaborate with otherwise in-conflict organizations? What would qualify as success for this boundary-organization-mediated collaboration? And what factors contribute to that success or lack thereof?

1.3.3 Legal considerations in standard-setting How laws might impact or govern these informal standard-setting processes might at first seem obscure. If a technical standard-setting body can be little more than a mailing list, the occasional meeting room and freely available documents, what legal considerations would even apply?

1.3.3.1 Anti-trust Directly applicable to any system for coordination can be legal rules against the development of trusts or cartels. For example, dividing up a market to reduce competition and increase prices can be a coordinated action that hurts consumers with higher prices and fewer new entrants, whether the conversation is formal or informal.

In the United States, antitrust law has historically been guided by the principle of consumer welfare, as laid out by Robert Bork (1978). That is, applications of

the Sherman Act are guided by whether consumers are hurt by the potentially anti-competitive behavior, through usurious prices or decreased choices. How “consumer welfare” is specifically defined, and whether “consumer welfare” alone is the appropriate way to analyze anti-trust enforcement, are openly debated questions. Anti-trust concerns and evaluations of consumer welfare have arisen around privacy in technical standards, as discussed in the findings.³⁶

Standard-setting bodies use transparency of decision-making and a documented system of due process as wards against trouble with antitrust enforcers ([Federal Trade Commission, Bureau of Consumer Protection 1983](#)). While substantive analysis requires substantial expertise, evaluating a “reasonable basis” for resulting standards, whether the standards are more pro- or anti-competitive, along with those basic procedural requirements has been the FTC’s approach to evaluating standard-setting organizations for antitrust ([Anton and Yao 1995–1996](#)). Internet standards organizations avoid making decisions or publishing documents that are specific to some set of vendors; this is considered a sound technical practice in general, but also helps to avoid legal entanglements for the participants. And “open” standards, where the resulting documents are made public, freely available and where participation in the process is generally open to anyone who wants to participate and where standards are adopted voluntarily, avoid many possible antitrust concerns, including the creation of cartels who could prevent new market entrants ([Lemley 1995–1996](#)).³⁷ Technical standards for interoperability instead can have an important pro-competitive purpose: the presence of a standard may inhibit the otherwise natural tendency towards market standardization,³⁸ where users of a networked technology like the Internet might flock to a single, proprietary offering, and enable competition between different vendors who implement the interoperable standard ([Lemley 1995–1996](#)).

1.3.3.2 Intellectual property Perhaps most common in technical standard-setting bodies, and especially in Internet standards groups, is some policy consideration for intellectual property rights in standards development. Most significant is patent licensing, but copyright and trademark can also play a role in organizational rules.

³⁶See [Competition and standard-setting](#).

³⁷However, Anton and Yao argue that “interface standards” (standards for interoperable communication, as in the case of Internet protocols) may be voluntarily adopted but still have anticompetitive effects because adoption while voluntary can still effectively be necessary for operation in a heavily networked marketplace ([1995–1996](#)).

³⁸As discussed above, this is an overloading of the term to encompass deeply conflicting concepts.

Patent licensing is of particular importance to standard-setting bodies because of the risk of “hold-up” (Contreras 2017). In brief, setting a standard can be a resource intensive process and once the standard is agreed upon, there can be substantial investment by implementers that depends upon that standard. If, after the standard is developed, a single player can assert a patent on some piece of the protocol design or the only feasible way to implement that standard (a “standards-essential patent”), then that player can extract an onerous rent on the implementers, requiring them to pay high licensing fees or face starting over on an entirely new standards design. This could discourage anyone from participating at all: why invest time and money in this collaborative process if it might be undermined after the fact? Worse yet, researchers document cases where firms apply for patents and actively manipulate standard-setting processes just to extract money from their competitors in patent licenses (Contreras 2017). In response, standard-setting bodies have set rules requiring disclosure of known patents or enforcing certain licensing terms.

Some have argued (Teece and Sherry 2002) that the preference for royalty-free licensing requirements or the patent licensing requirements in general might themselves be unfair to patent holders; that standard-setting participation can be a de facto requirement for wide adoption of a technology and that patent licensing rules in standard-setting bodies can be a cabalistic way of avoiding high prices. It is an argument that can be difficult for an engineer to assess with an open mind: that there would be something wrong with implementation-focused firms choosing to avoid patent fees or restrictions when designing new technology would seem very strange. Perhaps the more practical scope of the argument is that antitrust law should not altogether prohibit enforcement of standards-essential patents, as was feared in government moves that seemed to require reasonable and non-discriminatory licensing even when the patent holder was not a standards participant.

And standard-setting bodies have typically responded with their own organizational requirements rather than relying on regulatory imposition of some standard of fairness. W3C, for example, requires members to disclose patents they’re aware of and mandates royalty-free licensing of patents held by participants in a particular Working Group, with limited exceptions. Other standards groups accept licensing with royalty fees as long as the terms are fair, reasonable and non-discriminatory (FRAND).³⁹

³⁹For example, see the principles set out in the [OpenStand](#) group, including IETF, IEEE and W3C.

Copyright licensing of standards publications has also been a topic of debate in the Web standards community. While Web and Internet standards are freely and publicly available (that is, copyright is not necessary or useful for restricting access to the standards or extracting payment to read them), standard-setting bodies like W3C have often retained copyright on the published documents, largely for the purpose of preventing potentially confusing alternative versions to a canonical standard. Whether copyright is necessary or well-tailored to that use is unclear; many open source projects have approached trademark as an alternative. More permissive document licenses have surely not resolved all conflicts between WHATWG and W3C partisans.⁴⁰

1.3.3.3 Substantive law While legal considerations can affect procedural aspects of standard-setting laws specific to some sector can also influence standard-setting in that area.

For example, there is interest in using Web standards on user preferences and consent (the tools that make up DNT) to address implementation of the European General Data Protection Regulation (GDPR) or previous data protection directives. Technical standards are also practically necessary for realizing data portability requirements. Laws on accessibility of information and services to people with disabilities can create a market need for accessibility standards or even cite specific standards as required, “safe harbor”⁴¹ or example implementation.

While not directly in the area of Internet and Web standards, laws may incorporate standards by reference, in safety areas, for example. This may outsource or delegate regulatory decision-making to the private entities that set standards in those areas.⁴² It’s also been an area of critique where compliance with the law requires adhering to standards which may not be freely available to the public: Carl Malamud and Public.Resource.Org⁴³ in particular have fought legal cases (around copyright, in particular) in freely publishing standards referenced by public safety and building codes.

But while these laws and regulations can provide an incentive for developing or adopting standards in a particular area, those substantive rules have a different

⁴⁰See, for example, accusations of plagiarism when permissively licensed documents are copied, modified and republished.

⁴¹A safe harbor is one way to comply with a rule that is specifically acknowledged as satisfactory, removing further scrutiny or ambiguity.

⁴²For more discussion of legitimacy of delegated regulation, see [Drawing comparisons](#) below.

⁴³<https://law.resource.org/>

character of effect on the process of technical standard-setting than the legal considerations in anti-trust and intellectual property have.

1.3.3.4 Motivations for organizational policies In both areas of antitrust and intellectual property, standard-setting bodies – even consensus-based consortia – use organizational policies in response to potential exogenous legal constraints that might inhibit participation by individuals or firms. Standard-setting participation already has the disincentive of “free riding” – that freely available or “open” standards can be as easily used without the investment of time and resources into their development. If using a developed standard would incur the risk of patent infringement suits or the cost of patent licensing or if the informal standards meetings would prompt antitrust scrutiny, those dangers would minimize the potential economic gains of interoperability. Standard-setting groups are, as a result, responsive to these legal considerations that could create such disincentives.

1.4 Comparing governance models

Technical standard-setting is an important part of Internet governance but it’s often mistakenly analogized to legislating for the Internet. While standard-setting is a key point of coordination and implemented standards have profound impacts on design and use of the Internet, voluntary standards and consensus processes have a different force and character from legislation. Similarly, there may be some analogies to administrative law – rule-making and other regulatory authorities – but attending meetings and proposing new protocols is far from asserting power over how the Internet is used. As noted in the documentation provided to newcomers to the IETF:

If your interest in the IETF is because you want to be part of the overseers, you may be badly disappointed by the IETF. (“[The Tao of IETF: A Novice’s Guide to the Internet Engineering Task Force](#)” 2018)

Nevertheless, Internet governance, and technical standard-setting more specifically, can be a model for governance with the potential for collaboration that we should empirically evaluate.

1.4.1 What is Internet Governance The process of typing `nytimes.com` into your favorite Web browser’s address bar, hitting return and getting back the digital front page of that specific newspaper involves, when you interrogate the technical

details, an extraordinarily large number of steps. This exercise can be valuable pedagogy, in my experience, and it's also a famous interview question.⁴⁴

Many of those steps, many of the questions that make that discussion interesting, come down to determining how you the visitor can get an authoritative response – how you get *the* New York Times web page, how you're directed to web servers owned and operated by *the* New York Times, and there isn't confusion about who responds to what. The name `nytimes.com` has to be, in order to make the Internet work the way we have come to expect, universally registered to refer to that particular entity. When the domain name is translated into an Internet Protocol (IP) address – at the time of this writing, `151.101.1.164` – that address must refer to a specific server (or set of servers), it can't be in use by any other parties. The Internet (and it is capitalized in large part for this reason) requires a singular allocation of these resources, the names and numbers. At one time, that allocation was managed by a single person, specifically Jon Postel, and the recording of the allocation was done in a paper notebook.⁴⁵ As this became logistically infeasible (and later, when it became politically unacceptable), recording of names, numbers and protocol parameters was formalized as the Internet Assigned Numbers Authority (IANA) and by the late 1990s the IANA function was handled by a US non-profit corporation designed for that purpose, the Internet Corporation for Assigned Names and Numbers (ICANN).

The distribution of these resources can be complex and controversial. Regarding domain names, for example, a few questions arise:

- who gets what domain name,
- for how long,
- what if the domain name includes a registered trademark,
- who resolves disputes over a domain name,
- what if a domain name is being used for a criminal enterprise,
- what information should be made available about who owns a particular domain name,
- what top-level domains should there be,
- who gets to determine new ones,
- and on, and on.

⁴⁴<https://github.com/alex/what-happens-when>

⁴⁵Some sources refer to scraps of notebook paper, others refer to a notebook, but note it as “according to legend” (“History of the Internet” n.d.). In at least one interview (“Interview with Jon Postel” 1996), Postel refers to getting “the notebook” although it's not entirely clear if that's for the list of host addresses or the list of RFCs.

While the assignment of numbers might seem more straightforward, the exhaustion of the IPv4 space makes the job more challenging, and Regional Internet Registries (RIRs) subdivide the IP address space efficiently between large Internet service providers and users.

More obscure, the IANA function also includes maintenance of registries of protocol parameters,⁴⁶ values created or used by Internet standards where interoperability benefits from universal public registry. Port numbers were an early such case and a long registry of port numbers and services are still maintained.⁴⁷ It's useful to have a common convention that TCP connections used for accessing a Web server were made at port 80, and for different services to use different ports.

But while organizations exist to satisfy this allocation and registration of limited Internet resources, the standard-setting process enables the design of the protocols that use these resources. Protocols for identifying computers on the Internet, sending data between them, communicating the information necessary for efficient routing between networks, operating applications (email, the Web) on top of the Internet, securing Internet communications from eavesdropping or tampering – all these require standardized protocols, typically developed at the IETF, W3C or another standard-setting body.

And even with those standards developed and critical Internet resources allocated, the Internet depends upon relationships between individuals and organizations to keep communications flowing. Inter-domain routing, implemented through protocols (most specifically, BGP) developed in early days of the Internet when close relationships made security seem less necessary, still relies on trust developed between individuals at peer organizations. Mathew and Cheshire document that the personal relationships between larger network operators, developed over time through meetings and other interactions, and maintained through backchannel communications and resolving routing problems, make up an essential, decentralized part of maintaining orderly operation of the Internet (2010).

All these activities make up Internet governance,⁴⁸ a distinctive multistake-

⁴⁶<https://www.iana.org/protocols>

⁴⁷<https://www.iana.org/assignments/service-names-port-numbers>

⁴⁸“Internet governance” can either be narrowly defined as dividing up shared resources (IP address allocation and DNS name disputes) or broadly defined as the various activities (names and numbers, standards, peering agreements, trust relationships (Mathew 2014), etc.) for keeping the bits flowing. Or taking “governance” more broadly still, it can also refer to any government regulations related to the Internet, or to private actor actions that involve technical or self-regulatory implications for generation and distribution of content. There is no single accepted term.

“Internet governance” here is the distinctive set of activities that enables the definition and

holder model of decision-making that has maintained the operation of the Internet and the World Wide Web. Without these ongoing decisions, allocations and maintained relationships, the Internet would not function as the thing we recognize.

Multistakeholderism is a popular claim and a commonly-cited goal for Internet governance. In contrast to multilateralism (decision-making by sovereign governments, by treaty for example), multistakeholder processes are desired for not falling prey to ownership by a single government or bloc of governments and for responding to the interests of various kinds of groups, including business and civil society.

As part of a movement for “new governance,” the Obama administration called for multistakeholder processes as a responsive, informed and innovative alternative to government legislation or administrative rule-making (“[Commercial Data Privacy and Innovation in the Internet Economy: A Dynamic Policy Framework](#)” 2010; “[Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy](#)” 2012). Multistakeholder processes have also been suggested as alternatives during more recent drafting of potential federal privacy legislation. It is an especially relevant time to consider the lessons to be learned from Internet governance and from multistakeholder processes and to compare consensus-based technical standard-setting to other forms of governance.

1.4.2 Alternative governance models There is a hope for “collaborative governance” to promote problem-solving rather than prescriptive rule-setting. Freeman sets out five criteria for a collaborative governance rule-making process in the administrative law context ([Freeman 1997](#)):

1. problem-solving orientation;
2. participation by affected stakeholders throughout the process;
3. provisional conclusions, subject to further revision;

operation of the Internet, especially the allocation of resources, the development of interoperable protocols and the institutional or informal relationships that constitute its continued operation. Various forms of regulation (including all of law, norms, architecture, markets) that affect the Internet – laws to influence online commerce, the rules of large commercial platforms that govern use/speech of services, the technical designs of large Internet-enabled platforms – are fascinating, important, and not Internet governance, rather, simply that, governance that affects the Internet. Scholars interested in different governance debates that impact the use and development of the Internet will often look at that even broader scope; for example, Laura DeNardis and *The Global War for Internet Governance* ([DeNardis 2014](#)).

4. novel accountability measures;
5. an engaged administrative agency.

But the terminology of collaborative governance is used more broadly, and in some cases can push beyond even traditional public sector or government agency decision-making. On the broader side, Emerson et al. (2011) define collaborative governance as:

the processes and structures of public policy decision making and management that engage people constructively across the boundaries of public agencies, levels of government, and/or the public, private and civic spheres in order to carry out a public purpose that could not otherwise be accomplished.

It is this broader sense that fits the idiosyncratic nature of Internet governance in its different forms. And the model of collaborative governance regimes (CGRs) can provide the terminology (and some normative propositions or hypotheses) to describe the similarities and differences between public sector collaborative governance proposals and the techno-policy standard setting that my subsequent empirical work explores.

1.4.2.1 Regulatory negotiation Freeman evaluates regulatory negotiation (“reg-neg”) processes in the environmental health and workplace safety settings along the criteria for collaborative governance and finds them “promising” but with open questions regarding legitimacy and the “pathologies of interest representation.”

In a negotiated rule-making, a public agency starts a consensus-finding discussion with various stakeholders, and agrees (either in advance or after the fact) to promulgate rules under their legislatively-granted administrative authority that match that negotiated outcome. This kind of process is designed to decrease legal disputes over rules by involving as many of those parties in the negotiation itself (Harter 1982–1983) and to promote innovative problem-solving rather than adversarial interactions. In the case of regulating chemical leaks from equipment, the negotiation process that was expected to be a compromise on certain numbers and percentages of leaks turned into development of a new quality-control-inspired system, by both environmentalists and industry, that allowed “skipping” inspections when they were consistently positive and “quality improvement plans” when problems were discovered (Freeman 1997). In the case of EPA regulation of residential woodstoves, negotiation from states, environmentalists and the manufacturing

industry came up with an agreement on phased in rules with standardized labeling for the sale of new woodstoves where all of those parties agreed to defend the negotiated agreement in court (Funk 1987–1988).

Proponents identify the acceptance and stability of negotiated rule-makings (Harter 1982–1983) and the potential innovation in less adversarial settings (Freeman 1997). Critics of reg-neg oppose a negotiation process as an improper replacement of the administrative agency’s own expert determination of the public interest. That opposition can be on legal grounds – that the negotiated conclusion of the involved parties might go beyond or otherwise not match the particular legislative intent, an issue perhaps especially likely to happen with processes that look for novel re-framing of problems – or normative grounds – that negotiation between some group of parties will involve compromises or incomplete representation of stakeholders in a way that doesn’t adequately approximate the best interests of the public as a whole (Funk 1987–1988).

One open question that Freeman emphasizes is how these practices might apply in different contexts, and this study explores addressing user privacy concerns on the Web through multistakeholder standard-setting. There are certainly reasons to see several of those five criteria in the Internet standard-setting process.

Developing new protocols to enable new technology frequently lends itself to a problem-solving outlook (1) and the implementation and interoperability focus of Internet standards keep participants in that pragmatic mindset. Participants throughout the process include implementers, who remain involved throughout (2) design and deployment. While standards can be persistent in practice,⁴⁹ these “Requests for Comment” are expected to be revised regularly (3). Accountability is frequently considered in protocol design, with various measures including technical enforcement, market pressures, certification systems and governmental regulation. Perhaps least applicable in the analogy is the engaged government agency (5); while government representatives can and do participate in these consensus standard-setting fora, they are rarely a convener or among the most engaged. And while the literature of reg-neg suggests government agency rule-making authority as a kind of backstop to ensure legitimacy, resolution and support for the public interest in the negotiation, voluntary consensus standard-setting has, as we will see, no such direct governmental forcing function.

⁴⁹Hence one motivation for this project, the important and persistent infrastructural role that these protocol design decisions can play. Consider the anecdote commonly cited by Vint Cerf, that IPv4 was just to be a temporary prototype before the development of a production system.

1.4.2.2 Environmental conflict resolution Environmental conflict resolution (ECR) processes also represent a collaborative model for governance. This terminology also has different applications and meanings, but key properties of an environmental conflict resolution process seem to be: face-to-face meetings among a diverse group of stakeholders who have competing interests regarding some environmental outcome typically tied to a particular geographic location using some consensus-type process for determining a resolution, often (but not always) with the help of a neutral facilitator or mediator (Dukes 2004). The dispute might be dividing up the costs of cleaning up a spill or determining a plan for managing a set of natural resources.

ECR has been frequently practiced in the United States, providing a research corpus for evaluation. That research has included study of what are the appropriate success criteria to use in evaluating an ECR process and, what factors are connected to those success criteria. While not all participants in a process agree on whether it was successful, success can be measured in terms of: 1) whether agreement was reached, 2) what the quality of the agreement was and 3) how relationships between the participants improved. And more specifically, the quality of an agreement includes: a) how durably an agreement addresses key issues, b) the implementability of an agreement, c) the flexibility of an agreement to respond over time and d) the accountability of an agreement through monitoring or other compliance measures (Emerson et al. 2009, summarizing a broader set of research on ECR). Through multi-level analysis, Emerson et al. draw some conclusions on which beginning factors contribute to successful environmental conflict resolution, but emphasize that the intermediary step is effective engaged participation (2009). The change in working relationship stands out here because it isn't limited to the particular conflict or the particular agreement. Some scholars even identify the improvement in working relationships between parties as more important than the agreement over the initial conflict itself (Dukes 2004)!

1.4.3 Drawing comparisons Motivated by this work on collaborative governance and conflict resolution, I have tried to explore with my research participants their views on success criteria, including specifically the changes to working relationships. How well do the factors associated with successful conflict resolution explain the outcomes in technical standard-setting when it comes to policy-relevant challenges?

The success criteria and contributory factors in environmental conflict resolution have considerable overlap with Freeman's criteria for collaborative governance

problem-solving. Both cover pragmatism, participation, flexibility and accountability.

At the same time, we should identify factors of procedural and substantive legitimacy, as raised above. To the extent that government agencies rely on multi-stakeholder standard-setting processes to address disputes over public policy, there is a danger of regulatory delegation that may be unaccountable (Bamberger 2006), or put another way, that either regulatory agencies will be ‘laundering’ policy through a standard-setting process or they will be abdicating their responsibility to the public (Froomkin 2000; as cited by Boyle 2000). To this point, I have asked research participants about the fairness of the process and the fairness or quality of its outcomes.

1.5 The future of multistakeholderism for tech policy

We previously laid out a research agenda (Doty and Mulligan 2013), building on the suggestions of Waz and Weiser (2012) in a way specific to the development of techno-policy standards underway to address privacy issues on the Web. What are the impacts of multistakeholder processes in general, and multistakeholder techno-policy standards-setting processes in particular, on resolving public policy disputes for the Internet? How can we establish relative success and failure and what conditions affect those outcomes?

That agenda remains as relevant as ever in providing policy and policymaking advice given the interest in new governance and multistakeholder models. Privacy and security remain significant values of interest for this kind of approach and are of particular import with the Internet and the Web⁵⁰ but the set of public policy values where some collaborative, technical, problem-solving approach is desired only grows: harassment, abuse and free expression; diversity and representation; among others.⁵¹

This work places a downpayment on that research agenda. We can learn, I believe, from the history and practice of consensus standard-setting for the Internet and the Web and experiences of how it’s been used on matters of privacy and security. Nevertheless, this work also raises new questions on how and when technical standard-setting can be an effective multistakeholder process for tech policy issues.

⁵⁰See [Privacy and Security: Values for the Internet](#).

⁵¹See [Directions](#).

References

- “A Brief History of the Internet Advisory / Activities / Architecture Board.” n.d. Internet Architecture Board. Accessed September 8, 2018. <https://www.iab.org/about/history/>.
- Abbate, Janet. 2000. *Inventing the Internet*. MIT Press.
- Alpert, Jesse, and Nisan Hajaj. 2008. “We Knew the Web Was Big...” *Official Google Blog* (blog). July 25, 2008. <https://googleblog.blogspot.com/2008/07/we-knew-web-was-big.html>.
- Anton, James J., and Dennis A. Yao. 1995–1996. “Standard-Setting Consortia, Antitrust, and High-Technology Industries.” *Antitrust Law Journal* 64: 247. <https://heinonline.org/HOL/Page?handle=hein.journals/antil64&id=257&div=&collection=>.
- Bamberger, Kenneth A. 2006. “Regulation as Delegation: Private Firms, Decisionmaking, and Accountability in the Administrative State.” *Duke LJ* 56: 377.
- Benkler, Yochai. 2002. “Coase’s Penguin, or, Linux and ‘The Nature of the Firm’” *The Yale Law Journal* 112 (3): 369–446. <https://doi.org/10.2307/1562247>.
- Bennett, Colin J. 2010. *The Privacy Advocates: Resisting the Spread of Surveillance*. MIT Press.
- Berners-Lee, Tim. 2004. “How It All Started.” 2004. <https://www.w3.org/2004/Talks/w3c10-HowItAllStarted/>.
- Bork, Robert H. 1978. *The Antitrust Paradox*. Basic books New York.
- Boyle, James. 2000. “A Nondelegation Doctrine for the Digital Age.” *Duke LJ* 50: 5.
- Braden, R. 1989. “Requirements for Internet Hosts - Communication Layers.” RFC 1122. Request for Comments. RFC Editor. <https://tools.ietf.org/html/rfc1122>.
- Bradner, Scott. 1997. “Key Words for Use in RFCs to Indicate Requirement Levels.” RFC 2119. Request for Comments. RFC Editor. <https://tools.ietf.org/html/rfc2119>.
- Brooks, Sean, Michael Garcia, Naomi Lefkowitz, Suzanne Lightman, and Ellen Nadeau. 2017. “An Introduction to Privacy Engineering and Risk Management in Federal Systems.” NIST Internal or Interagency Report (NISTIR) 8062. National Institute of Standards and Technology. <https://doi.org/https://doi.org/10.6028/NIST.IR.8062>.
- Bruant, David. 2013. “The W3c Is a Restaurant.” *Long-Term Laziness* (blog). October 8, 2013. <https://longtermlaziness.wordpress.com/2013/10/08/the-w3>

[c-is-a-restaurant/](#).

- Cargill, Carl F. 1989. *Information Technology Standardization: Theory, Process, and Organizations*. Newton, MA, USA: Digital Press.
- Cerf, Vinton, and Robert Kahn. 1974. "A Protocol for Packet Network Intercommunication." *IEEE Transactions on Communications* 22 (5): 637–48.
- "Commercial Data Privacy and Innovation in the Internet Economy: A Dynamic Policy Framework." 2010. National Telecommunications and Information Administration, Internet Policy Task Force. <https://www.ntia.doc.gov/report/2010/commercial-data-privacy-and-innovation-internet-economy-dynamic-policy-framework>.
- "Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy." 2012. White House. <http://www.whitehouse.gov/the-press-office/2012/02/23/fact-sheet-plan-protect-privacy-internet-age-adopting-consumer-privacy-b>.
- Contreras, Jorge L. 2017. "Technical Standards, Standards-Setting Organizations and Intellectual Property: A Survey of the Literature (With an Emphasis on Empirical Approaches)." SSRN Scholarly Paper ID 2900540. Rochester, NY: Social Science Research Network. <https://papers.ssrn.com/abstract=2900540>.
- Crocker, Stephen D. 2009. "How the Internet Got Its Rules." *The New York Times*, April 6, 2009, sec. Opinion. <https://www.nytimes.com/2009/04/07/opinion/07crocker.html>.
- DeNardis, Laura. 2009. *Protocol Politics: The Globalization of Internet Governance*. MIT Press.
- . 2014. *The Global War for Internet Governance*. Yale University Press.
- Dessart, George. n.d. "Encyclopedia of Television - Standards and Practices." The Museum of Broadcast Communications. Accessed August 30, 2018. <http://www.museum.tv/eotv/standardsand.htm>.
- DiMaggio, P J, and W W Powell. 1983. "The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields." *American Sociological Review* 48 (2): 147–60.
- DiMaggio, Paul. 1982. "The Structure of Organizational Fields: An Analytical Approach and Policy Implications." In *SUNY-Albany Conference on Organizational Theory and Public Policy*.
- Doty, Nick, and Deirdre K. Mulligan. 2013. "Internet Multistakeholder Processes and Techno-Policy Standards: Initial Reflections on Privacy at the World Wide Web Consortium." *Journal on Telecommunications and High Technology Law* 11. http://www.jthtl.org/content/articles/V11I1/JTHTLv11i1_MulliganDoty.PDF.

- Dukes, E. Franklin. 2004. "What We Know about Environmental Conflict Resolution: An Analysis Based on Research." *Conflict Resolution Quarterly* 22 (1-2): 191–220. <https://doi.org/10.1002/crq.98>.
- Emerson, Kirk, Tina Nabatchi, and Stephen Balogh. 2011. "An Integrative Framework for Collaborative Governance." *Journal of Public Administration Research and Theory*, May. <https://doi.org/10.1093/jopart/mur011>.
- Emerson, Kirk, Patricia J Orr, Dale L Keyes, and Katherine M Mcknight. 2009. "Environmental conflict resolution: Evaluating performance outcomes and contributing factors." *Conflict Resolution Quarterly* 27 (1): 27–64. <https://doi.org/10.1002/crq.247>.
- fantasai, and Florian Rivoal. 2020. "W3c Process Document." World Wide Web Consortium. <https://www.w3.org/2020/Process-20200915/>.
- "FAQ — WHATWG." n.d. Web Hypertext Application Technology Working Group (WHATWG). Accessed August 23, 2018. <https://whatwg.org/faq>.
- Federal Trade Commission, Bureau of Consumer Protection. 1983. "Standards and Certification: Final Staff Report." Washington, D.C. <https://catalog.hathitrust.org/Record/001535861>.
- Ford, Paul. 2014. "The Group That Rules the Web." *The New Yorker*, November 20, 2014. <https://www.newyorker.com/tech/elements/group-rules-web>.
- Freeman, Jody. 1997. "Collaborative Governance in the Administrative State." *UCLA L. Rev.* 45: 1.
- Froomkin, A. Michael. 2000. "Wrong Turn in Cyberspace: Using ICANN to Route Around the APA and the Constitution." *Duke LJ* 50: 17.
- . 2003. "Habermas@Discourse. Net: Toward a Critical Theory of Cyberspace." *Harvard Law Review* 116 (3): 749–873. <https://doi.org/10.2307/1342583>.
- Funk, William. 1987–1988. "When Smoke Gets in Your Eyes: Regulatory Negotiation and the Public Interest - EPA's Woodstove Standards." *Environmental Law* 18: 55–98.
- Gillies, James, and R Cailliau. 2000. *How the Web was born: the story of the World Wide Web*. Oxford: Oxford University Press.
- Guston, David H. 2001. "Boundary Organizations in Environmental Policy and Science: An Introduction." *Science, Technology, & Human Values* 26 (4): 399–408. <http://www.jstor.org/stable/690161>.
- Haberman, Brian, Joseph Lorenzo Hall, and Jason Livingood. 2020. "Structure of the IETF Administrative Support Activity, Version 2.0." RFC 8711. Request for Comments. RFC Editor. <https://rfc-editor.org/rfc/rfc8711.txt>.
- Harter, Philip J. 1982–1983. "Negotiating Regulations: A Cure for Malaise." *Georgetown Law Journal* 71: 1. <https://heinonline.org/HOL/Page?handle=hein.journ>

- als/glj71&id=17&div=&collection=.
- “History of the Internet.” n.d. APNIC. Accessed August 10, 2018. <https://www.apnic.net/about-apnic/organization/history-of-apnic/history-of-the-internet/>.
- “Interview with Jon Postel.” 1996. January 29, 1996. <http://oceanpark.com/papers/postel.html>.
- “ISO 1806:2002 - Fishing Nets -- Determination of Mesh Breaking Force of Netting.” 2002. December 2002. <https://www.iso.org/standard/28360.html>.
- “ISO/IEC 27001 Information Security Management.” 2013. Information Security Management Systems. <http://www.iso.org/cms/render/live/en/sites/isoorg/home/standards/popular-standards/isoiec-27001-information-securit.html>.
- Jacobs, Ians. 2009. “Frequently Asked Questions (FAQ) about ISOC and W3c.” World Wide Web Consortium. December 2009. <https://www.w3.org/2009/11/isoc-w3c-faq>.
- Kelty, Christopher M. 2008. *Two Bits: The Cultural Significance of Free Software*. Duke University Press. <https://twobits.net/>.
- Kesteren, Anne van. 2017. “Further Working Mode Changes.” *The WHATWG Blog* (blog). December 11, 2017. <https://blog.whatwg.org/working-mode-changes>.
- Latour, Bruno. 1987. *Science in Action : How to Follow Scientists and Engineers Through Society*. Harvard University Press. <http://www.amazon.com/dp/0674792904>.
- Leiner, Barry M., Vinton G. Cerf, David D. Clark, Robert E. Kahn, Leonard Kleinrock, Daniel C. Lynch, Jon Postel, Larry G. Roberts, and Stephen Wolff. 2009. “A Brief History of the Internet.” *SIGCOMM Comput. Commun. Rev.* 39 (5): 22–31. <https://doi.org/10.1145/1629607.1629613>.
- Lemley, Mark A. 1995–1996. “Antitrust and the Internet Standardization Problem.” *Connecticut Law Review* 28: 1041–94. <https://heinonline.org/HOL/P?h=hein.journals/conlr28&i=1051>.
- Lind, Edgar Allan, and Tom R. Tyler. 1988. *The Social Psychology of Procedural Justice*. Springer.
- Maathuis, I., and W. A. Smit. 2003. “The Battle Between Standards: TCP/IP Vs OSI Victory Through Path Dependency or by Quality?” In *ESSDERC 2003. Proceedings of the 33rd European Solid-State Device Research - ESSDERC '03 (IEEE Cat. No. 03ex704)*, 161–76. <https://doi.org/10.1109/SIIT.2003.1251205>.
- Mathew, Ashwin. 2014. “Where in the World Is the Internet? Locating Political Power in Internet Infrastructure.” <https://www.ischool.berkeley.edu/research/publications/2014/where-world-internet-locating-political-power-internet-infrastructure>.

- Mathew, Ashwin, and Coye Cheshire. 2010. "The New Cartographers: Trust and Social Order Within the Internet Infrastructure." SSRN Scholarly Paper ID 1988216. Rochester, NY: Social Science Research Network. <https://papers.ssrn.com/abstract=1988216>.
- "Memorandum of Understanding Between W3c and WHATWG." 2019. May 28, 2019. <https://www.w3.org/2019/04/WHATWG-W3C-MOU.html>.
- Moon, Sangwhan, Travis Leithead, Arron Eicholz, Steve Faulkner, and Alex Danilo. 2017. "HTML 5.2." W3C Recommendation. World Wide Web Consortium. <https://www.w3.org/TR/2017/REC-html52-20171214/>.
- O'Mahony, Siobhán, and Beth A. Bechky. 2008. "Boundary Organizations: Enabling Collaboration Among Unexpected Allies." *Administrative Science Quarterly* 53 (3): 422–59. <https://doi.org/10.2189/asqu.53.3.422>.
- Ostrom, Elinor. 2015. *Governing the Commons*. Cambridge university press.
- Overell, Paul, and Dave Crocker. 2008. "Augmented BNF for Syntax Specifications: ABNF." STD 68. Network Working Group. <https://tools.ietf.org/html/rfc5234>.
- Polletta, Francesca. 2004. *Freedom Is an Endless Meeting: Democracy in American Social Movements*. University of Chicago Press. <http://books.google.com/books?id=snuq08KeC2EC>.
- Postel, Jon. 1981a. "Internet Protocol." 791. Request for Comments. RFC Editor. <https://rfc-editor.org/rfc/rfc791.txt>.
- . 1981b. "Transmission Control Protocol." 793. Request for Comments. RFC Editor. <https://rfc-editor.org/rfc/rfc793.txt>.
- Rockefeller, John D. 2013. *A Status Update on the Development of Voluntary Do-Not-Track Standards*. <https://www.commerce.senate.gov/public/index.cfm/2013/4/a-status-update-on-the-development-of-voluntary-do-not-track-standards>.
- Saxenian, AnnaLee. 1996. *Regional Advantage*. Harvard University Press.
- Simcoe, Timothy. 2014. "Governing the Anticommons: Institutional Design for Standard-Setting Organizations." *Innovation Policy and the Economy* 14 (January): 99–128. <https://doi.org/10.1086/674022>.
- Star, Susan Leigh, and James R. Griesemer. 1989. "Institutional Ecology, 'Translations' and Boundary Objects: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-39." *Social Studies of Science* 19 (3): 387–420. <https://doi.org/10.1177/030631289019003001>.
- Teece, David J, and Edward F Sherry. 2002. "Standards Setting and Antitrust." *Minn. L. Rev.* 87: 1913.
- "The Art of Consensus: A Guidebook for W3c Group Chairs, Team Contact and Participants." n.d. World Wide Web Consortium. Accessed August 25, 2018.

- <https://w3c.github.io/Guide/>.
- “The Tao of IETF: A Novice’s Guide to the Internet Engineering Task Force.” 2018. IETF. November 8, 2018. <https://www.ietf.org/about/participate/tao/>.
- Tyler, Tom, and David Markell. 2010. “The Public Regulation of Land-Use Decisions: Criteria for Evaluating Alternative Procedures.” *Journal of Empirical Legal Studies* 7 (3): 538–73.
- Waz, Joe, and Phil Weiser. 2012. “Internet Governance: The Role of Multistakeholder Organizations.” <http://www.silicon-flatirons.org/documents/publications/report/InternetGovernanceRoleofMSHOrgs.pdf>.
- “We’re ISO: We Develop and Publish International Standards.” n.d. International Organization for Standardization. Accessed August 30, 2018. <https://www.iso.org/standards.html>.
- “Web IDL.” 2018. <https://heycam.github.io/webidl/>.
- Werle, R., and E. J. Iversen. 2006. “Promoting Legitimacy in Technical Standardization.” *Science, Technology & Innovation Studies* 2 (1): 19–39. <http://www.sti-studies.de/articles/2006-01/werle.htm>.