

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Enacting Privacy in Internet Standards

Permalink

<https://escholarship.org/uc/item/3h04r0m8>

Author

Doty, Nicholas

Publication Date

2020

Peer reviewed|Thesis/dissertation

Enacting Privacy in Internet Standards

By

Nicholas P Doty

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Information Management & Systems

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Deirdre K. Mulligan, Chair

Professor Kenneth A. Bamberger

Professor Jenna Burrell

Fall 2020

Abstract

Enacting Privacy in Internet Standards

by

Nicholas P Doty

Doctor of Philosophy in Information Management & Systems

University of California, Berkeley

Professor Deirdre K. Mulligan, Chair

The functionality of the Internet and the Web are determined in large part by the design of technical standards that allow for interoperable implementations. Those design decisions are important both in terms of functionality and in maintaining basic public policy values including accessibility, freedom of expression, privacy and security. This is one instance of a phenomenon variously referred to as *values-in-design* or from another angle *technological delegation*: fundamental matters of public policy importance can be determined or regulated by software architecture, in much the way that urban architecture has.

Technical standard-setting bodies, like the Internet Engineering Task Force (IETF) and the World Wide Web Consortium (W₃C), are multistakeholder organizations that host groups that define Internet standards. In contrast to multilateral bodies (like the United Nations) or state bureaucracies or firms, participants are drawn from competing companies across different industry sectors, as well as governments and civil society groups, and make decisions based on rough consensus and working implementations. This represents a distinct form of *new governance* where collaborative development of technical and policy solutions can be enacted outside traditional legislative or administrative bodies.

Standard-setting participants are most often engineers; their discussions are technical and wide-ranging, distributed between in-person meetings, email lists, online chat and version-controlled text and software. Web standards include HTML, the foundational markup language that defines Web pages, and Do Not Track, a syntax and system for communicating user privacy preferences about online behavioral tracking. Some standards, like HTML, have been widely adopted and support some of the world's most used software, while others, like DNT, have seen limited adoption and little direct effect.

This work seeks to understand how privacy and other values get enacted in the technical standards and running software that make up the Internet we use. At the highest level, it considers these two related research questions:

1. What are the impacts of multistakeholder techno-policy standards-setting processes on resolving public policy disputes for the Internet?
2. How do the views of the designers of the Internet's underlying protocols affect the privacy of Internet and Web users?

To start to answer these, I have explored the community of standard-setting participants and their beliefs about privacy and security in their lives and work. And I have investigated the unusual consensus decision-making process used in technical standard-setting and perspectives, from newcomers and long-time participants, on its fairness and efficacy and how it applies to values such as privacy.

Private interviews with participants working on Do Not Track and other standards provided candid and diverse perspectives on the concept of privacy, multi-stakeholder processes and the role of technical standards for interoperability and for policy-related areas. And the extensive documentation and technologically-mediated communication methods of these Internet standard-setting venues enabled some supplementary quantitative analysis of the patterns of participation.

Multistakeholder standard-setting processes bring together diverse participants from a wide variety of organizations with a wide variety of backgrounds and goals. Individuals navigate a tricky balance of being both experts collaborating and representatives negotiating. While this novel cross-boundary process provides real opportunity, it also provides real difficulties of bad faith behavior, entrenchment and conflict. And while access and transparency of processes may improve upon some alternatives, technical standard-setting continues to be Western-oriented, male-dominated and intensely time-consuming.

Participants hold competing views of privacy and recognize that the priorities and concerns of Internet users may vary widely. While privileged community members may not have as much to risk when it comes to their own online privacy, the distinct wants and needs of their own children provide a compelling touchstone. Sometimes the work of privacy has been reduced to compliance with laws or best practices, but it remains an area where professionals actively pursue debate and development of policy.

The social, legal and technical architecture of the Internet and the Web determine so much about the lives of people around the world and deserve the attention of research on their impacts. Using consensus-based multistakeholder processes

focused on interoperability for the Internet presents real, new opportunities to enact privacy by intentionally taking advantage of handoffs between social, technical and organizational factors. At the same time, this work highlights some of the challenges to convening for, equitably designing, agreeing on and implementing these techno-policy standards. Training people in intersecting disciplines, developing systematic processes and building technical and decisional tools can all contribute to better support for privacy, security and the other fundamental but still contested values we want to see in the Internet.

Acknowledgments

Thanks to my advisors, colleagues, friends and family.

Thanks to colleagues and friends in the Doctoral Research and Theory Workshop and Berkeley writing groups over the last several years, both for sharing their own work so that I could learn from their practices and for reviewing several chapters of this dissertation and providing thorough, kind and constructive feedback.

Thanks to colleagues at the World Wide Web Consortium (W3C) and Internet Engineering Task Force (IETF) and in the tech field generally for encouraging me to learn and contribute. While I fear that these settings are not as accessible or welcoming as they could or should be, I know that I have benefited from being able to participate.

Thanks to my interviewees for their trust in me, candor about their own experiences and generosity with their time and expertise.

Thanks to my dissertation committee – Deirdre K. Mulligan, Jenna Burrell and Kenneth A. Bamberger – for their patience and insight; and to Paul Duguid for his useful advice, regarding my dissertation prospectus and in general.

Thanks to faculty and students at UC Berkeley and the School of Information for introducing me to ideas and challenging the ones I had; in particular: Calvin Morrill for his seminar on the sociology of law and organizations; Jenna Burrell for her introduction to qualitative research methods; the Classics reading group for critical analysis of the history of technology; and Technology & Delegation for exploring many of the fundamental ideas of this work.

Thanks to collaborators with whom I've had the pleasure to work on the following projects (incomplete lists, forgive me):

- BigBang: Seb Benthall, Niels ten Oever and Corinne Cath;
- privacy design patterns: Mohit Gupta and Jaap-Henk Hoepman;
- essentially-contested concepts: Deirdre K. Mulligan and Colin Koopman;
- handoffs: Deirdre K. Mulligan, Helen Nissenbaum and Richmond Wong; and,
- the Center for Technology, Society & Policy: Galen Panger, Deirdre K. Mulligan, Steve Weber and Anno Saxenian.

Thanks to the staff of the School of Information for keeping things running, in particular: Kevin Heard and Gary Lum for providing technical infrastructure and Inessa Gelfenboym Lee for helping me through writing and bureaucracy.

Thanks to John MacFarlane for pandoc and for standardizing CommonMark; thanks to Kieran Healy for LaTeX templates and for popularizing plain text social science; and thanks to the many contributors who have built the SciPy ecosystem.

Thanks to conveners and participants at the venues where I have had the opportunity to present work in progress, get essential feedback and learn from others, including: the International Workshop on Privacy Engineering, the Computing Community Consortium's Privacy by Design Workshops, the Trustbusters Workshop at the Symposium on Usable Privacy and Security, TPRC, the Web 2.0 Security and Privacy Workshop, South By Southwest, the TRUST Center, the Human Rights Protocol Considerations Research Group, the Alan Turing Institute's Protocol Governance Group, and, most especially, the Privacy Law Scholars Conference.

Funding

Funding to support the research included in this dissertation and related work has been provided by:

- the Berkeley Center for Law & Technology, through a grant from the Nokia Corporation;
- the US Department of Homeland Security (DHS) under Grant Award no. 2006-CS-001-000001, under the auspices of the Institute for Information Infrastructure Protection (I3P) research program managed by Dartmouth College;
- the Team for Research in Ubiquitous Secure Technology (TRUST) Center, which receives support from the National Science Foundation (#0424422);
- INSPIRE: Value-Function Handoffs in Human-Machine Compositions, awarded by the National Science Foundation (#1650589);
- the Center for Technology, Society and Policy (CTSP), which received major funding from Google;
- the International Computer Science Institute (ICSI) Science of Privacy Labet, supported by the National Security Agency (NSA);
- the University of California at Berkeley and the School of Information, through fellowships and travel grants.

Contents

0	Introduction	1
1	Internet Standard-Setting and Multistakeholder Governance . . .	4
1.1	What is a standard	5
1.2	The consensus standard-setting model	8
1.3	Organizational structure	23
1.4	Comparing governance models	33
1.5	The future of multistakeholderism for tech policy	40
2	The Ethics of Engineering	41
2.1	Engineering is inherently ethically-laden	41
2.2	Separation vs. integration	42
2.3	Ethics in organizations, professions and individuals	45
2.4	Engineering impacts for values of the Internet	49
3	Privacy and Security: Values for the Internet	51
3.1	Definitions and contentions	52
3.2	Relationship between privacy and security	54
3.3	Cases in this work	57
3.4	Encrypting the Web, a “handoff”	58
3.5	Do Not Track, a “handoff”	68
4	A Mixed-Methods Study of Internet Standard-Setting	86
4.1	Setting the scene	87
4.2	Studying standard-setting at different scales	91
4.3	Researcher position	93
4.4	Interviewing	96
4.5	Analyzing communications	103
4.6	Ethics	106
5	Findings	109
5.1	Connecting research questions to themes	109
5.2	How standard-setting accommodates, succeeds and fails	111

5.3	Competition and standard-setting	134
5.4	Individuals vs organizations in standard-setting process .	144
5.5	Who participates and why it matters	154
5.6	How participants see privacy	170
5.7	Towards integration	189
6	Directions for future work	192
6.1	A triad for interventions	192
6.2	Recognizing future handoffs	195
	References	197
	Appendices	212
	Appendix: Interview Guide	212
	Appendix: Privacy-as-x	215

This dissertation is also readable in Web and PDF form at:

<https://npdoty.name/enacting-privacy/>

Corrections, comments and links to related or responding work will be maintained there.

o Introduction

What are technical standards and why look so closely at them? In [Chapter 1: Internet Standard-Setting and Multistakeholder Governance](#), I provide background on the consensus standard-setting model and how standards are developed for the Internet and the Web. I then consider how Internet governance and multi-stakeholder standard-setting models compare to calls for new and collaborative governance approaches and set out the first high-level research question for this project: what are the impacts of multistakeholder techno-policy standards-setting processes on resolving public policy disputes for the Internet?

In [Chapter 2: The Ethics of Engineering](#), I build a philosophical argument for engineering as an inherently ethically-laden practice and trace the competing impulses for separating out and more deeply integrating ethical considerations into technical design. Given the ethical importance of engineers and engineering, I introduce the second research question for this project: how do the designers of the Internet's underlying protocols view privacy and how do their views ultimately affect the privacy of Internet and Web users?

Why look at privacy? In [Chapter 3: Privacy and Security: Values for the Internet](#), I explain why privacy and the related but distinct property of security are and have been values of particular importance and ongoing contestation in the design of the Internet and the Web. To illustrate, I describe two cases where there has been a handoff of responsibility for some conception of privacy between technical, legal, organizational and individual actors. First, the movement to encrypt the Web, deploying security technology to maintain user privacy from network surveillance and intrusion; and second, Do Not Track, an effort to develop a cooperative mechanism to enable user choices about privacy from online behavioral tracking.

Having set out the theoretical lens, the key questions and the focus of my research, [Chapter 4: A Mixed-Methods Study of Internet Standard-Setting](#) describes the mix of methods I used to study the distributed, mediated, networked setting that is Internet standard-setting. Different methods can be most useful at different scales. This project involves interviews with standard-setting participants, with sampling across a distinctive set of dimensions, to elicit their individual and personal perceptions and feelings, as well as expertise, about privacy and about the working process of standardization. And at a macro scale, I use quantitative analysis of mailing list archives to measure the demographics of, and social connections between, participants. I have focused my empirical inquiry on Do Not Track and the related standardization process where I was most deeply involved.

[Chapter 5: Findings](#) lays out my findings from qualitative interviews and

quantitative social network analysis that speak to those two research questions: how multistakeholder techno-policy standard-setting process affects public policy values and how participants' views of privacy affect the privacy of Internet users.

For my first research question, I explore themes related to the process itself, the stakeholders involved, the roles of individuals and organizations and the patterns of participation. I review the standard-setting process itself at multiple stages and how the process either fails or succeeds at accommodating what participants saw as a mix of good faith and bad faith behavior and a range of diverse perspectives and backgrounds. Regarding a particular debate over anti-trust, I show the different purposes that transparency has in standard-setting processes and in governance generally that influence decision-making in the moment and how it's understood and interpreted later. Transparency also influences the role of policymakers who participate in multistakeholder process, making it difficult to effectively apply a soft touch when discussions happen in private. I describe the tradition of individual participation in Internet standard-setting and the complicated interactions that arise from competing views of the individual's role as an expert in a largely technocratic collaborative process or as a representative of a stakeholder group in a political process of balancing policy views. Because representation often affects how we see legitimacy, I provide some demographic metrics on who is participating, including initial results on gender disproportion, and detail the differing views of how many sides are involved, which may influence entrenchment and how to identify opportunities for cooperation. And in considering what makes standard-setting succeed, I document the importance of formal and informal leadership and the dense community structure of overlapping groups of repeat participants.

For my second research question, I look in detail at what participants in technical standard-setting processes related to privacy think about privacy itself: what their conceptions of privacy are and what privacy concerns they identify for themselves and in their work. Conceptions of privacy vary widely, from confidentiality to autonomy to freedom from intrusion, but most importantly participants explicitly anticipate and respond to the variety of views and priorities they expect from users. Interviewees identify some kinds of data as especially sensitive because of the potential for chilling effects, inferences about intimate areas of life or the risk of very direct intrusions. I also describe how they understand and are motivated by the privacy interests of others, from their own children to Internet users at large.

What this leaves for the future is the question, or rather, the challenge, of what practices we could use in technical standard-setting to more effectively enact

privacy for the Internet and the Web. Having now characterized the process and the participants; having described the contestations over the concept of privacy and the purpose of standard-setting processes; and having identified some of the difficulties in using multistakeholder process to resolve these debates, in [Chapter 6: Directions](#), I describe a triad of areas for intervention: people, processes and tools. Looking forward, we can recognize and analyze potential handoffs of responsibility between people, laws and technology and develop novel collaborative solutions to enact privacy, security and other human values.

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,
- 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

¹h/t Richmond Wong

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 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

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.



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.

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 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. 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).

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

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. 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

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).

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.

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 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

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

Layers of the Internet

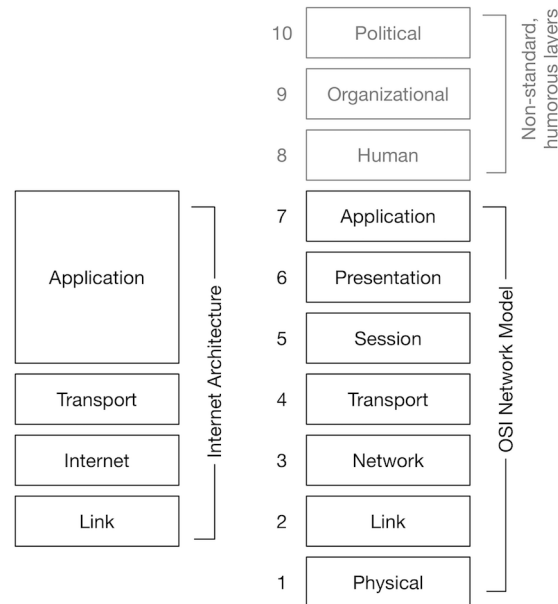


Figure 2: Layers of the Internet, both the OSI seven-layer model and the TCP/IP four-layer model (Braden 1989), aligned.

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 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

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](#).

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.

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 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)

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”.

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 (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, as you might expect, a particular representation of technology-focused firms.¹⁵

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).

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>.

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:

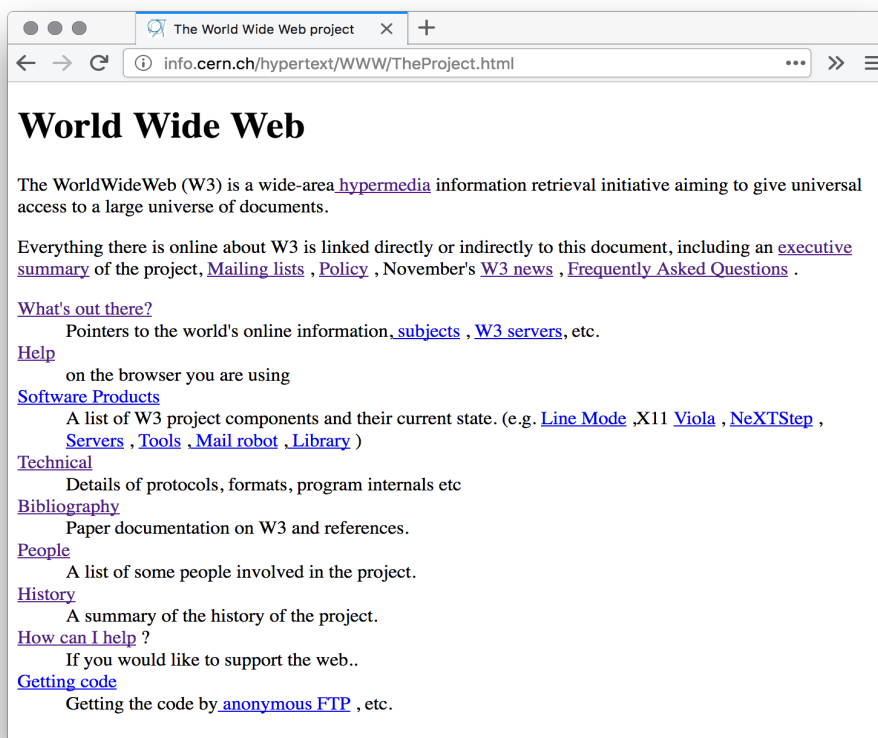


Figure 3: The first ever Web site is again operational on CERN's servers, with early descriptions of the Web, its operation and motivations.

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

<http://www.internetpolicystreams.com/>.

(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 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

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.

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?"

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.

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,

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).

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 kilogram), 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

²¹<https://fidoalliance.org/>

²²<https://kantarainitiative.org/>

²³<https://www.oasis-open.org>

²⁴<https://www.nist.gov/>

²⁵<https://www.iso.org>

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.

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

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](#).

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

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.

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 confusion 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

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

³⁰ h/t Daniel Griffin, for the lovely term

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

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).

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

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.

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;
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.

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

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.

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 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

³⁵See [Individuals vs organizations in standard-setting process](#).

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) – 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

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

(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.

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

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.

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

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 permis-

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

sive 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 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”

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/>

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

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

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.

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.

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.

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 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,

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 ([J. Freeman 1997](#)):

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

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](#)).

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 (J. 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 (J.

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.

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

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.

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*.

2 The Ethics of Engineering

2.1 Engineering is inherently ethically-laden

In studying the ethical implications of the Internet and the Web (or indeed of technology in general), one might reasonably ask: why study the engineers at all? Why not just study the users of technology, or the business incentives for tech-focused corporations, or the specific details of software artifacts? Scholars of science, technology and society do examine all those things, with different research focuses,⁵² but I see a strong philosophical basis for exploring the ethics and ethos of those who engineer and design technology. While arguments against a view of tools as purely neutral are widespread and diverse, I am guided by the view described by José Ortega y Gasset of technology as a particularly human act.

Ortega's argument is more essentialist than other arguments for the ethical implications of technology. Technology is not ethical just because it has a particular set of consequences and those consequences happen to be ethical ones; rather, technology is a set of choices about the good life.

To follow the argument step by step (Ortega y Gasset and Miller 1962):

1. technology is the distinctly human act of changing or reforming nature;
2. it is characteristic of man to employ technology, "the adaptation of the medium to the individual" (p. 96), to address her necessities;
3. but technology is not limited to creating biological necessities;
4. and indeed man seems to consider those "superfluous" things to be essential to life: "Not being, but well-being, is the fundamental necessity of man, the necessity of necessities" (p. 99); to sum it up:

Man, technology, well-being are, in the last instance, synonymous.
(p. 100)

Ortega concludes from the synonymy that the direction of technology is inherently subjective, as a result of the different views of the good life:

Whereas life in the biological sense is a fixed entity defined for each species once and for all life, life in the human sense of good life is

⁵²The designer-artifact-user spectrum is only one way of classifying researchers or the object of research in technology and society, but it can be an interesting one for information science colleagues. South Hall whiteboards have explored this in spectrum and triangle form.

always mobile and infinitely variable. And with it, because they are a function of it, vary human necessities; and since technology is a system of actions called forth and directed by these necessities, it likewise is of Protean nature and ever changing. (p. 101)

If you accept the subjectivity of our desires and what makes for a good life and you accept Ortega's argument that technology just is the reformation of nature to bring about those various superfluous aims, then, he argues, you should not accept that there is a singular progression of technology. To do so would be to "assume that man's vital desires are always the same and that the only thing that varies in the course of time is the progressive advance towards their fulfillment. But this is as wrong as wrong can be" (p. 102).

The synonymity of technology and well-being and the potential losses of technology that comes from different desires or changed circumstances leads to the argument that engineering ought to be conceived of broadly, rather than as a narrow, neutral, technical item. Ortega is arguing for a holistic view of technology and the good life.

para ser ingeniero, no basta con ser ingeniero⁵³

Engineers have to be more than just engineers because their work is the work of constructing human well-being and that view of well-being may change: "the social, economic and political conditions under which he works are changing rapidly" (p. 104).

2.2 Separation vs. integration

There are two fundamentally competing impulses over the role of the engineer and the engineering process in the ethical implications of a system. One is towards separation. It's considered sound engineering practice to maintain a "separation of concerns": the efficiency, modularity, reusability and testability of code all benefit from making each component self-contained and focused on its own task. An analogous philosophy argues for that separation in the process of developing new technology; the engineer focuses on the mechanism, not the policy ("[Mechanism Not Policy](#)" 2005), on the how, rather than the what. A developer might say that

⁵³*En:* to be an engineer, it is not enough to be an engineer. I first encountered this quote in [Morozov \(2013\)](#).

the choice for how the system is supposed to work is “above my pay grade” or that the quality of a piece of code is determined by whether it meets the specification provided by the customer or the manager. Engineers may choose to “punt” a decision to be resolved later or elsewhere, either for pragmatic concerns or to enable flexibility or choice by some other party (Doty 2015b). Architects of the Internet have recommended a principle of accommodating “tussle” of different priorities of different stakeholders, including by “designing for choice” by different parties in a communication, because conflict is inevitable and unresolvable (David D. Clark et al. 2002).⁵⁴

Even the holistic view of technology as ethics can include this perspective: at times even Ortega cites engineers as a rank below “poets, philosophers [and] politicians” because the engineer is dependent on their analysis of the values of human life. Or to use the cogent example of the development of the atomic bomb, Richard Sennett introduces us to the argument over the engineer’s ethical culpability and involvement (2008), positioning Hannah Arendt as arguing for the subservience of the engineer (1958).

A counter-acting impulse is toward integration of ethical concerns into the development process. Scholars and practitioners both have argued that technical decisions are not “pure,” “apolitical” or “neutral.” There are infamous examples of choices of technical architecture with profound, concrete and durable impacts on basic questions of public policy, like the height of overpasses and the inaccessibility of parks to people without wealth (Caro 1975).⁵⁵ These cases of *technological delegation* emphasize the impracticality of a separation approach.⁵⁶ At times, scientists and engineers have spoken up to express their strong ethical perspectives, bolstered by their knowledge and participation in the development of influential

⁵⁴Clark et al. explicitly consider “mechanism, not policy” and describe it as “too simplistic” but still a valuable principle in trying to separate out pieces of the system more or less likely to involve tussles between parties.

⁵⁵However, note that while Caro’s example of overpasses to prevent public transit access is illustrative and well-documented, most of *The Power Broker* portrays Moses not as the skilled engineer (indeed, he has no engineering training), but as a skilled legislative aide, manager and architect of public opinion.

⁵⁶This conversation can very easily get confused when we talk about the attribution of values. Some get upset when Latour writes about technical agents and accuse him of a basic fallacy of attributing mental states and intentions to inanimate objects. When critics speak of the bias or politics of algorithms (Tufekci 2016; Winner 1980), some technical audiences are confused because the algorithm itself has no apparent political content or skewed intent. It is often the *choice* of algorithm that has a political impact and the decision of the designer of a robotic agent that carries a moral weight.

technologies; to continue with the atomic bomb case, Einstein co-authored a post-war manifesto (1955), directed towards politicians and government leaders and arguing for pacifism.

This impulse towards intentional integration prompted the creation of a proto-field of academic study in “values in design” (VID): a community of interdisciplinary scholars recommending consideration of ethics and human values in the design of technology and infrastructure, rather than waiting for those implications to be seen and addressed after the fact (Knobel and Bowker 2011). But recognizing that values considerations can be relevant to technical design decisions does not automatically make it easy to integrate them (Flanagan, Howe, and Nissenbaum 2008):

It is one thing to subscribe, generally, to these ideals [either the ideals of liberty, justice, enlightenment, etc. or the responsibility to take them into account], even to make a pragmatic commitment to them, but putting them into practice, which can be considered a form of political or moral activism, in the design of technical systems is not straightforward.

With that inherent integration recognized, there have been prominent attempts in global politics to use the design process proactively to buttress values of interest.⁵⁷ Regulators increasingly call for “privacy-by-design”⁵⁸ with the hope that software built to support privacy will have fewer of the unanticipated and troubling breaches of privacy during its use. Privacy by design may include: default settings for more private modes; data minimization so that technical systems collect or retain only the granularity of information needed for a particular purpose; and, audits and organizational controls to limit misuse of personal data.

With perhaps less political attention, similar reasoning has been used to promote a security development lifecycle (Lipner 2004) and to consider other system properties (internationalization, accessibility, performance) throughout software development, and in each case we can see the struggles to enact those values and system properties, struggles including at least epistemological (what really is the value in each case) and practical (what methods and practices are best to bring it about) barriers (Flanagan, Howe, and Nissenbaum 2008).

⁵⁷Flanagan, Howe and Nissenbaum called this the “pragmatic turn” to “values as a design aspiration” (2008).

⁵⁸Cavoukian popularized the term and devised one specific process, but the approach in more general terms has been adopted by policymakers and software firms around the world.

It might be taken as obvious or simply accepted that the design of new communications technology has impacts on important public policy values and that there are ethical implications to the design decisions that engineers make. But these impulses in tension explain why that widely-accepted importance does not translate straightforwardly into how most ethically to design technology. Our study must recognize these competing principles and engineering practices. The ethics of engineering will include both accommodating diverse, conflicting uses and embedding some fundamental values.

2.3 Ethics in organizations, professions and individuals

Questions of how architectural decisions with ethical implications are made are often answered with high-level explanations based on economic incentives or legal constraints. (“Why does Google track my online browsing activity?” “Because tracking provides a higher return-on-investment in online advertising and Alphabet Inc. is a for-profit shareholder-value-maximizing firm.”) Market dynamics are no doubt important in the direction of technology firms and economic explanations will be useful in explaining corporate actions. But in this work I will primarily seek to examine the backgrounds, motivations and decisions of individuals (including software engineers and other participants in technical standard-setting) and the dynamics of working groups and professional communities.

I believe economic arguments do not have the explanatory power or richness that other social scientific analyses can provide and that free-market economics alone cannot account for the relevant architectural decisions made by engineers and others in the development of the Internet and the Web. This belief is informed by my understanding of:

- the process of software engineering – architecture is made up of many small-scale detailed decisions made by those who are intimately engaged with the material; and,
- the design of the Internet – the Internet is effected collaboratively and cooperatively.

That particular architectural decisions by individual engineers have meaningful ethical consequences is also informed by the philosophical arguments of the previous sections. If we accept the holistic view, as Ortega argues, that technology

doesn't just have ethical implications but is by its nature a defining of what is a good life and if we accept that integration of values into the design of technical systems is at least sometimes preferable, then we should, as researchers, look at the perspectives and practices of individuals engaged in engineering and design to more fully understand these ethical-technical decisions.

2.3.1 An ethos of engineering To understand the ethical practices and commitments of this Internet engineering community, it is useful also to consider the *ethos*: their character or guiding concepts. Coleman (2012) describes the interplay of hacker ethics and aesthetics. While she is careful to point out that there is no singular hacker ethic, Coleman identifies political strains of liberalism (free speech, inalienable labor) connected to the deep satisfaction (*eudaimonia*, even) of tinkering and subversion of systems within F/OSS contributors.

Software engineering shares with other types of engineering an impulse to “build,” “make” or “create.” That impulse can develop an ethic to do something, to build something in part exactly because one can do so (Doty 2013). To solve a difficult problem, even without a particularly remunerative or societally valuable outcome is often considered sufficiently motivating reward. A common method of recruiting software engineers is extolling the set of “hard problems” to work on. We can see both exploratory motivations (a la climbing a mountain “because it’s there”) and a motivating sense of independence (showing that you can do it on your own, through use of technology) here.

At the same time, technology faces a challenge in response: just because you can do something, should you?

Or, almost conversely, given the privilege of those few who can make potentially great differences through the creation and use of technology, are engineers doing their best to live up to that opportunity? This became a pointed question in the responses to the suicide of Aaron Swartz (aaronsw) and in local debates about tech company social responsibility, displacement and housing in San Francisco. And it spawned many recitations and riffs on the opening from “Howl” (Ginsberg 1955), applied to the apparent lack of ambition or importance of software development among Web giants:

The best minds of my generation are thinking about how to make people click ads.

While this earliest version of the quote is from Jeff Hammerbacher (formerly of Facebook) in 2011 (Vance 2011), it became a common, even blasé, criticism



Figure 4: “your scientists were so preoccupied with whether or not they could, they didn’t stop to think if they should.” A quote from the movie *Jurassic Park*, now commonly used as a meme to humorously indicate that some novelty is foolish or irresponsible. For example: “Hey Jim Comey, listen to Jeff!”

of priorities in software development.⁵⁹ It’s not clear how precise the analogy is, whether Hammerbacher or others intended a reference to drug addiction, homosexuality or the artist as an outcast in materialist society as described by the Beat poets, or if “I saw the best minds of my generation” is simply a memorable introduction that can somewhat ironically be used to describe well-educated, ambitious computer programmers. It is consistent, though, in suggesting that the tech industry and individual software engineers make substantial impacts and in lamenting the loss of an opportunity to apply that intellectual energy to some higher goal. The ethos of capability and impact is tied to an ethical aspiration.

2.3.2 Professional ethics Ethical norms spread through a profession, as well as through organizational hierarchies or personal social ties. Famously, the Hippocratic Oath is used as a formal example of an ethical code in medicine, a shared common agreement that among other commitments doctors shall, first, do no

⁵⁹I mostly stopped keeping track after 2014, but even then, it was turning into a meta-joke because it was so widespread: <https://pinboard.in/u:npdoty/t:bestminds/>.

harm. In traditional engineering (civil and mechanical, in particular), a similar moral commitment is present in the Ritual of the Calling of an Engineer (written by Rudyard Kipling) or in the oath of the Order of the Engineer (“Obligation” 2018):

As an engineer, I pledge to practice Integrity and Fair Dealing, Tolerance, and Respect, and to uphold devotion to the standards and dignity of my profession, conscious always that my skill carries with it the obligation to serve humanity by making best use of the Earth’s precious wealth.

Professionalization can be a way for an obligation to the public to be maintained, even when it might be contrary to the particular interests of an individual or firm. The sociology of law has shown evidence that the professional background and training of lawyers can distribute norms across national boundaries (Carruthers and Halliday 2006). Software engineering may not have a code and professional societies with the same pervasiveness as in medicine, law or engineering, but ethical codes, ethical education and professional organizations⁶⁰ are present and it’s clear that professionals in engineering and technology are asking the same questions as other professions about their commitments to society.

Calls for a Hippocratic Oath or more rigorous ethical codes of conduct for practitioners in software engineering and data science are widespread. That might be an indication that the existing codes of professional ethics are not widely known. But criticism of engineering ethics codes and their utility or focus is long-standing.

For example, Luegenbiehl and Puka note (1983) the historical basis of ethical codes in engineering as driven by an interest in professionalization, criticizes them (unfairly, I would say) for not being exhaustive guides on ethical conduct, and notes the individualism (perhaps inherited from legal and medical ethics) that may not be appropriate for engineering practices that we know affects the wider public. Lynch and Kline (2000) argue for considering the ethics of everyday and non-technical parts of engineering practice, rather than focusing too narrowly on whistleblower moments and case studies of conflicts with amoral management. Davis (1991) argues for the utility of a code of ethics as part of a profession, in solving the coordination problem of individual, ethically-minded engineers overcoming a client’s or manager’s request. But he also concludes that engineers have

⁶⁰The Association for Computing Machinery is currently revising its code of ethics. More specific organizations develop training and certification for sub-fields; IAPP for privacy, for example.

this professional and ethical motivation not because of any familiarity with the text of a code of ethics, but because it's part of "thinking like an engineer."

How codes conceive of their obligations can provide an explanation of (or indicate the presence of) a cultural perspective towards the profession. For example, some codes will focus on an obligation to the public while others may emphasize the responsibility to a particular client, with likely different results in professional attitudes. Stark and Hoffmann identify different motivating metaphors in ethical codes that represent different professions (or different parts of a broader computing or data science profession) and correspond to responding to different values and prioritizing different constituencies (2019). Professional codes can contribute to credibility or to benevolence or both, and computer engineering has unfortunately not had a focus on benevolence. They quote Kate Crawford in noting, "data ethics needs to ask, 'what kind of world do we want to live in?'" Indeed, if we see engineering as technical work inherently and explicitly asking that question (as Ortega suggests) – how will tools shape a different world and which different world do we want? – then we can see engineering ethics as not just professional behavior or appropriate stakeholder harm-reduction, but an essential aspect of engineering.

In each of these reviews, the ethical impact of engineering and the construction or focus of an ethical code for the engineering profession has emphasized the distinctive practice of engineering itself, whether it's "thinking like an engineer" or what makes up the day-to-day practice of engineering work. The ethos of engineering itself is difficult to specify but it meaningfully impacts how ethical practices are approached by individuals and communicated across organizations and professions.

2.4 Engineering impacts for values of the Internet

This chapter has reviewed the inherent ethical impacts of engineering. Given the outsized role that Internet engineering and the choices of many individual software engineers have for values such as privacy, this research seeks to understand how privacy is or is not supported by those who develop the Internet and the Web. We must ask, how do the designers of the Internet's underlying protocols view privacy as a value? And how do their views ultimately affect the privacy of Internet and Web users?⁶¹

For this research project, I focus on public policy areas with a special valence for the Internet and the Web: security and privacy. Why those values, and how

⁶¹These make up dissertation Research Question 2.

responsibility for those values is allocated and distributed, is explained in the following chapter: [Privacy and Security: Values for the Internet](#).

3 Privacy and Security: Values for the Internet

Privacy and security are just two values among many that can be enacted within a technical design. Accessibility, accountability, archivability, fairness, free expression, internationalization, justice, neutrality, performance and many other values can all be affected by the particular technical architecture.⁶² However, privacy and security have played an outsized role in the history and use of the Internet and the Web, despite the clichéd and largely inaccurate notion that the original design of the Internet ignored security. Because of the decentralized architecture of the Internet and the end-to-end property of its design, security is a challenge to achieve, while being a pre-requisite for the use of the Web for electronic commerce. Because the Internet is, most of all, an information medium that billions of people use to communicate, protecting privacy and control over the flow of personal information is a fundamental task, especially as users contribute more of their personal thoughts in increasingly popular social media applications.

We could imagine an alternative history of the World Wide Web that didn't prioritize these applications — ecommerce, personal communications, social media — one that was more limited to the accessible library of information originally imagined by Tim Berners-Lee. In theory, such a Web might see privacy or security as less fundamental issues. With fewer commercial applications, confidentiality, integrity and availability may have been less pressing properties for development; if the Web were more a reading platform than one where users generated content themselves, privacy issues, while germane, might be less inherently essential. Arguing for a necessary history of the Web from its origins to its current form is counter to good historiography; in this case, it is also unnecessary. There are reasons to support the notion that the success of the Internet and the Web made it likely that commercial applications would be developed and that without commercial applications the infrastructure would not be as substantial or as popular. As noted previously, from the earliest days of the Internet, email and personal communications were essential drivers of the infrastructure. Similarly, the architecture and history of the Web suggest that user-provided content of some form would be supported, whether through a “read-write Web”⁶³ or more centralized social

⁶²See the “values in design” concept and the trend towards integration, as described in [The Ethics of Engineering](#), previously.

⁶³A concept long-advocated by Tim Berners-Lee and popularized in the Read/Write Web blog ([MacManus 2003](#)), where users can contribute to web pages as easily as they browse them.

media. As interesting as these alternative likely histories are, the fact remains that ecommerce and social media have been large, popular, driving applications of the Web, and applications that are particularly likely to involve security and privacy issues. As such, it's fruitful to look at these values, even as we recognize that other values have also been important to the development of the Web and that different values may support different applications in the future.

To begin with, let's define, or at least scope, some of the basic terminology.

3.1 Definitions and contentions

“Security” can mean many different things to different people and in different cultural contexts. While some might immediately think of the locked doors of a bank vault (an access control view), others might think of the safety of basic needs. The Japanese word “anshin” is used in some contexts as a translation of security, but describes more broadly a sense of peace of mind, tied to confidence, familiarity and knowledge (Okumura, Shiraishi, and Iwata 2013).

In the fields of network or information security, what is considered the classical model defines security as a property of a system that satisfies three objectives: confidentiality, integrity and availability (the C-I-A triad).⁶⁴ While critiques and extensions to the confidentiality-integrity-availability model are common, most researchers in these fields continue to rely on something similar; this research uses “security” to refer to these objectives in computer/information security unless otherwise noted.

Contentions about the definition of “security” are mild in comparison to the myriad differences over “privacy.” A classical definition is that privacy is control over personal information, as presented by Alan Westin, considered a founder in the field (1967). That definition mirrors early definitions of security: an access-control approach based on satisfying a particular privacy or security policy. However, many scholars have noted limitations to this “informational privacy” definition; that it doesn't capture intrusions into our daily lives or substantively capture what is distinctive about violations of that control over information. Many practitioners rely on a concept of “fair information practices” or “fair information practice principles” (FIPPs), drawn from a Department of Health Education, and

⁶⁴The original source identifying these objectives as fundamental to security is unknown. An early reference identifying them as the most common goals of a security policy is a report from Dave Clark and David Wilson: “A Comparison of Commercial and Military Computer Security Policies” (1987). Notably, this is the same Dave Clark known for design of the Internet architecture.

Welfare report from 1973 (Department of Health, Education and Welfare 1973) and the Organization for Economic Cooperation and Development from 1980 (Organization for Economic Cooperation and Development 1980) and still very present in the Obama administration’s proposed Consumer Privacy Bill of Rights (“Consumer Data Privacy in a Networked World: A Framework for Protecting Privacy and Promoting Innovation in the Global Digital Economy” 2012). Philosopher Helen Nissenbaum argues for a theory of “contextual integrity” (2004) to explain our privacy views based on the flows of personal information: not absolute access control policies but instead expectations built up from social and legal norms.

You might prefer one privacy definition over another or find one more often used in a particular setting, but increasingly it seems clear that “privacy” is an *essentially-contested concept* for which we will not and should not settle on a single definition. Following the characteristics laid out by Gallie (1956), privacy is: *appraisive*, a valuable achievement; *complex*, with multiple dimensions including objectives and justifications; *open*, changing in salience over time in response to different technological and social circumstances; and, finally, subject to *progressive competition*, where ongoing debates over the concept can contribute to better understanding privacy (Deirdre K. Mulligan, Koopman, and Doty 2016). As a practical matter, this suggests conducting research to anticipate and uncover, rather than foreclose, different approaches to privacy. And as we note in that work, contestation of privacy has important implications for design:

- debates over a single definition of privacy will not be conclusive, and so it will be more useful to describe particular concepts or privacy goals in a particular context;⁶⁵
- a static list will not be able to anticipate all privacy concerns, and so designers can benefit both from looking very concretely at specific user needs or concerns and at a higher-level from understanding the object of privacy and identifying where it can be supported in the technical architecture; and,
- because contestation will continue, designers should anticipate and accept this openness.

⁶⁵See, for example, [this two-year discussion of a definition of privacy and whether it’s necessary or useful for IETF specification work on the ietf-privacy mailing list](#). In Do Not Track discussions, participants debated whether defining “privacy” was useful for scoping or an unimportant academic matter.

Instead of a singular definition of privacy, then, we end up with meta-analyses of privacy concepts. Dan Solove argues for a Wittgensteinian “family resemblance” approach and sets out a large taxonomy of different actions that might constitute privacy violations (2006). Colin Koopman and Deirdre Mulligan devise a privacy analytic to map out theories of privacy on a large number of dimensions, including the purpose of protecting privacy and from whom one is protecting their privacy (Koopman and Mulligan 2013; Deirdre K. Mulligan, Koopman, and Doty 2016).

By necessity, then, this research does not rely on a single, narrow definition of “privacy” for its inquiry. Further, as a methodological matter, foreclosing any dispute on the definition or sense of privacy might lead to missing that same dispute within the community or communities in question. How privacy is differently defined by the engineers and other participants in technical standard-setting is itself a research question. Existing work has looked at the models of privacy evident in the work of computer scientists working in security and privacy (Danezis and Gürses 2010) and in the nascent field of privacy engineering (Gürses and Alamo 2016).

Care is taken, as a matter of research method, not to “prime” or load a particular meaning of the term “privacy” during interviews with participants.⁶⁶ This method is more than contingently important, because one possibility to be explored is that, because of the openness in response to technological change of values particularly impacted by the Internet, what privacy is may not only be debated among engineers, but materially constructed by them.

For the purpose of scoping my own inquiry, I focus on privacy as the family of values related to norms and controls over flows of information about people and freedom from intrusions.

3.2 Relationship between privacy and security

Why consider privacy and security together? Aren’t these separate values that need to be distinguished in order to determine the distinctive effects and factors related to privacy?

There is some truth to the common cliché that “you can’t have privacy without security.” That is, systems that are vulnerable to attacks that break the properties of confidentiality or integrity typically can’t guarantee users control over how information about them is collected, used or disclosed. This is true in more than the most naive sense in which security is necessary for a system to provide any other

⁶⁶See interview guide in the appendix.

value — if a system is not available, then it cannot provide any of its functionality; if a system cannot provide integrity, then it could have been altered to counter some other value for which it was designed. For example, any conception of privacy that includes keeping some information secret or controlling access to a piece of information will be undermined by violations of confidentiality: if a system is vulnerable to threats where an attacker can access information she is not intended to be able to access, then the system is less likely to provide contextual integrity or effective controls over information disclosure.

From the perspective of Internet architecture, security may be more relevant at lower layers – e.g. establishing secure channels of application-agnostic communication – while privacy may be more significant at higher layers – e.g. user controls over information disclosure in particular applications.⁶⁷ Braman identifies privacy as a topic of concern from the earliest days of Internet architecture design as described in the first ten years of RFCs, with confidentiality and access control of particular importance for protecting information on hosts or transmitted through the network (2012).

In addition to security as a pre-requisite for (or lower layer to) privacy, there are also cases where privacy and security overlap. One common reason for conflating security and privacy is the assumption that privacy *just is* confidentiality. It's popular to claim that this conflation is simply erroneous; however, if we accept that privacy is plural and essentially contested, it's more difficult to flatly discount such a theory. What we can say is that most typical definitions or theories of privacy include protections beyond confidentiality. That is, privacy-as-confidentiality is an uncommonly *narrow* conception of privacy. That said, those same typical definitions (including both control over personal information and contextual integrity) would count many confidentiality violations as prototypical violations of privacy: many concepts of privacy *include* confidentiality.

For example, the National Institute of Standards and Technology (NIST), in seeking to improve and systemize the engineering practices for privacy, has drafted an evolving set of privacy engineering objectives, to serve a similar functional purpose to the C-I-A triad. In initial drafts, the list of objectives included confidentiality (as defined in related security engineering documentation) to explicitly mark an overlap between privacy and security (NIST 2014).⁶⁸

⁶⁷This may be a common perspective, but I'm not sure whether it's published or documented as a design principle. I attribute it to presentations by Alissa Cooper.

⁶⁸Partly in response to public comments, a subsequent draft finding uses "disassociability" instead, with a definition distinct from confidentiality, and more like "unlinkability" (NIST 2015).

Distinct from the layer model (privacy on top of security) described above, there might also be cases where a lack of privacy undermines security. Designs for security that rely on trust in participants might have a vulnerability if the personal privacy of an individual is compromised. For example, the confidentiality of classified information depends on the reliability and lack of coercion of those cleared to receive that information; if the intimate details of a person's life are accessed, a blackmailer may be able to obtain government secrets.⁶⁹ Similarly, some authentication mechanisms rely on limited flows of information about a person; if an attacker can unexpectedly easily determine your birthdate and addresses of previous residences, they may be able to impersonate you to your bank.

While there are substantive connections between privacy and security in the design of Internet protocols, an additional motivation to consider these values together is their integration in the *practice* of privacy and security engineering work. As later sections will demonstrate, the work of identifying and mitigating privacy concerns and security concerns share techniques (like threat modeling), expertise and people. Even if values can be, conceptually, separately defined and considered, if the engineering efforts are themselves combined, then understanding and improving the practice of privacy and security engineering requires exploring the values together.

As an empirical matter, efforts for coordinated security and privacy review have become more integrated in recent years. One explanation is that, in addition to the inherent connections between accepted security properties and common conceptions of privacy, the historical context of a changing political and technological atmosphere has shifted privacy to depend more deeply on traditional security objectives. That openness is a piece of privacy's essential contestedness. In a historical review of privacy, we can note how privacy (at least in Western society) was broadly conceived in the late 19th century as a freedom from harassment or publicity – a response to photography and newspapers; and in the mid-20th century as a concern about unfair or unaccountable analysis in newly-available large, computerized databases. I believe we can see a similar shift over shorter timeframes in the conception of online privacy in the 21st century. When a plethora of online tracking mechanisms and corresponding behavioral advertising companies appeared in the early 2000s, the privacy concern of protection from corporate profiling was heightened; after the Snowden revelations in 2013, a shift in effort and attention was made towards privacy from large-scale government surveillance and securing infrastructure. That a concern was heightened during a particular

⁶⁹h/t Daniel Griffin

time doesn't imply that it was absent otherwise; government surveillance was not a wholly new concern after 2013 and online corporate data collection remains a privacy issue (not just because the same infrastructure is relevant to government access). Similarly, the privacy torts about unwanted publicity didn't disappear after the 20th century. But these historical shifts and the competing concepts of privacy they highlight are, I argue, reflected in the work on engineering privacy on the Internet and its increasing integration with security.

While this work focuses on privacy and security as fundamental values in tension on the Internet, what we learn from the design for these values can inform, and be informed by, research on the design of other values. In particular, there is much to learn from experiences with accessibility and internationalization; and I hope this research can contribute to work on diversity and freedom from harassment.⁷⁰

3.3 Cases in this work

Following these shifts in the concept of privacy, let us look at two cases, with different conceptions of privacy and where there is a change, or potential change, in the distribution of responsibility for protecting privacy. In each, we can see a “handoff” in the larger socio-technical system and the manner of these shifts can help us uncover what value is being supported and how.

First, I look at the movement to [encrypt the Web](#), including designing, advocating for and deploying new security technology to maintain privacy from network surveillance and intrusion; and, second, I consider [Do Not Track](#), an effort to develop a cooperative user choice mechanism for protecting privacy from online behavioral tracking, which will be the topical focus of my empirical work.

⁷⁰See [Directions](#).

3.4 Encrypting the Web, a “handoff”

3.4.1 Cafe Confusion You sit down at the little cafe on the corner, much like the cafe where I’m sitting and writing this story. It’s a lazy afternoon, you order a latte and while you’re waiting for it, you open your laptop and connect to the only open network: FreeLittleCafeWiFi. Why not open your email and see if your sister wrote you back? And before you even get to Yahoo! Mail, you think, actually, you should check Facebook and see if there are any new pictures of the nephews, and who’s coming to that party on Saturday night.

There’s a handful of other people in the cafe: a teenager probably from the high school across the street, a man with glasses tapping away on the next great American novel, a woman working on a presentation. Everyone is using a laptop.⁷¹

After a few minutes scrolling, scrolling, scrolling through the news feed, you notice that there’s a new post: it says it’s from you, but you certainly didn’t write a post of fart jokes. The teenager across the room snickers as she closes up her laptop and leaves the cafe.

Frustrated, you delete that post. What else did they see in your account? Hadn’t you made sure to log in with the lock icon, and wasn’t that lock icon supposed to protect you? Could people in the same cafe always see what you were doing online? Who else could do that? Shouldn’t that be against the law? Shouldn’t people know better? Did you do something wrong? *Who or what was really responsible?*

3.4.2 Handoffs That uneasy question of responsibility arises from an unsettled combination of technical, legal and social processes. There is an implicit distribution among technical and non-technical means of assurance for a particular value of interest within a sociotechnical system. How you are able to communicate with friends and family over the Internet and whether those communications are secured from prying eyes and tampering depends both on the architecture of the Internet and the World Wide Web and on legal and normative protections of privacy and security.

In the example above, the mischievous teenager in the cafe might have used a small plugin called “Firesheep” that makes such eavesdropping and impersonation a straightforward point-and-click measure. The author of Firesheep, Eric Butler, makes his case for who is responsible: companies operating these websites *should have* implemented more widespread security with HTTPS for all connections and

⁷¹It’s 2010, say, and not quite everyone is using smartphones all the time yet.

creating Firesheep was a way to expose this problem more clearly.⁷² That position is well-argued, but the question of normative responsibility doesn't have singular answers; the Web could also be designed such that legal, rather than technical, protections were what disincentivized the attacker; or a technical system could prioritize accountability and auditing over protection against an attack. Or, as the argument is sometimes made from a certain reductionist perspective, the value of privacy might simply not exist in a certain setting, *caveat emptor*.⁷³

Collaborators have defined a **handoff** as the transfer of a function or the responsibility for some value from one actor to another between two different configurations of a system (Deirdre K. Mulligan and Nissenbaum 2020). A movement towards encrypting the Web, more specifically to broadly increasing the fraction of Web traffic communicated over TLS-encrypted channels, is such a handoff, shifting the responsibility for security from that unsteady combination of factors to a deployed technical system of HTTPS in browsers and servers.

In this case, I detail the different actors that make up the socio-technical system that is the Web, its diversity of goals, and the handoff to a set of technical guarantees for providing the value of security in online browsing. We can see how the distribution of responsibility has changed, both a general shift in paradigm and a particular triggering event. How that handoff is being negotiated and implemented shows how values can be conceived, debated and enacted in a complex, distributed system.

3.4.2.1 System overview The Web as a **socio-technical system** is complex in both its makeup and function.⁷⁴ Billions of end users use web browsers on personal smartphones, laptops, smart televisions, desktop computers at their local library or Internet cafe. Web sites that those users visit are produced by newspapers, governments, corporations, non-profits, individual hobbyists; those sites are hosted on servers ranging from tiny low-powered devices sitting on a bookshelf to enormous server farms with distributed locations around the world. Interconnection between those end users and those servers typically happens over the Internet, itself an even larger system; communications typically hop from a local WiFi network, to a commercial ISP, to some series of backbone providers, to a

⁷²The very directly titled presentation “Hey Web 2.0: Start protecting user privacy instead of pretending to.”

⁷³Perhaps, *caveat usus*, but I don't have the ideal translation for “user.”

⁷⁴A similar overview of the Web as a socio-technical system opens the Do Not Track handoff case as well – these are written so that they can be read individually as discussion drafts for the handoffs model.

commercial network provider, to a CDN or commercial server, and back again. Depending on routing protocols, peering arrangements, server distribution and local network infrastructure, those hops may cross many national boundaries and may take many different routes or all pass through a single undersea network cable.

A incomplete summary of relevant (human) **actors**:

- end users
- Web site developers
- browser developers
- ISP administrators
- advertising network executives

Or considering other *types of actors*, we might also identify key pieces of hardware or software: network switches, Web browsers, fiber-optic cables. Or institutional actors: diverse privacy laws in the US and EU, nation-state intelligence agencies, commercial security companies, organized crime syndicates.

3.4.2.2 Diversity of goals Given the diverse users of the system and the broad spectrum of actors that compose it, the Web also has a wide range of goals or functions. Many people use the Web for personal communications: checking their email accounts, posting messages to their social network accounts, reading and writing blog posts. Commerce is a common set of functions: companies provide services for sale; people buy both digital and physical products; online advertising is widespread; media companies provide entertainment services. In part because the Web and the Internet can be used for quick personal communication, intelligence agencies also use the network for surveillance of different kinds, to review the messages of particular targets, to map social networks based on communications metadata, to detect new security threats.

For this case study, we will look at a single goal, or a related set of functions for which the Web could be designed: securing the communications between people and services.

Security is a broad, multi-faceted concept: consider the C-I-A triad (D. D. Clark and Wilson 1987) and Japanese *anshin* (Okumura, Shiraishi, and Iwata 2013). We identify security as a **value**. Confidential and integral communications could be considered the relevant **goal**, or the goal might simply be communicating with others and the intended **constraint** is for those communications to be widely available while being free from tampering or eavesdropping.

3.4.2.3 Paradigmatic changes

3.4.2.3.1 Implicit trust in the network Securing communications on the Web could potentially be accomplished through many different configurations of the socio-technical system. Historically (this is an overgeneralization, but stay with us), Web traffic was typically not encrypted between the endpoints. In order to facilitate online commerce – as users were concerned about entering their credit card numbers into such a new and less-understood system – transport-layer security standards were developed and many sites implemented that security for specific security-focused operations, like entering payment information or sending passwords for logging in to accounts.

This **paradigm** – *occasional security with user confirmation* – presumes trust in a range of network intermediaries. Assurance of the confidentiality of communications with your email provider, say, depended on the discretion of the ISP and other network backbone providers. Integrity of communications against modification by intermediaries was simply not provided as a technical matter; occasionally network providers would insert advertising or notifications for the subscriber.⁷⁵ Laws, regulations, norms and market forces could provide an incentive for those network intermediaries towards securing unencrypted communications against unwanted disclosure or troublesome tampering. Because those companies were typically regulated and within the jurisdiction of national governments, law enforcement or intelligence agencies had the technical and legal capability, at their discretion, to intercept any particular Internet communication. Relying on norms and legal backing, network operators, technical designers or expert users may have expected that such discretionary activity would be abnormal in the United States or other liberal democracies. Technical enforcement (transport-layer security, based on a PKI of certificate authorities) was most often used for explicitly sensitive data. Technical protection against downgrade attacks was limited or absent. Implicitly or explicitly, users had the responsibility to confirm through browser UI that a connection was secure before entering credit card numbers, passwords or sensitive information in order to obtain that technical support of confidentiality. Understanding error messages about the security of connections was challenging and users are faced with various seals and lock iconography with unclear implications (Sunshine et al. 2009).

⁷⁵For example, Comcast has documented their notification system that inserts Javascript into web pages visited by the user (Chung et al. 2011).

3.4.2.3.2 Surveillance revelations as trigger To identify a singular trigger for the re-thinking and re-engineering of such a massive sociotechnical system, even limited to this particular function of secure communications, would be handwaving over a complex history. The position that the whole web should use HTTPS was common in certain communities before 2013, for various reasons related to privacy and security.

However, there are indications of a turning point in rhetoric and a substantial change in the momentum of action towards a new handoff configuration that can be related – in time and by explicitly-stated motivation – to the revelations in 2013 of widespread mass surveillance by the NSA and GCHQ.

Statements from engineers at the time indicate an acknowledgement of the previous handoff between state and technical actors, as well as the shift.⁷⁶

3.4.2.3.3 A new paradigm: encrypted transport everywhere Driven by evidence of tampering with web traffic by ISPs and other intermediaries and widespread passive surveillance by state actors, recommendations for Web security moved towards encrypted transport (HTTPS) being ubiquitous or expected for all (or most) kinds of Web usage. Rather than relying on the user to know when HTTPS was appropriate or necessary and manually confirm its use, servers were provided with the means to suggest or force usage of secure communications.⁷⁷ In this **paradigm** – *security for all Web traffic, driven by server and browser*, the user is out of the loop; Web communications are to have confidentiality, integrity and authentication by default, without user intervention, or even user understanding. In terms of threat modeling, the network is considered an attacker; widespread passive surveillance is directly addressed and not just for commercial activity but for personal information, various powerful Web capabilities, and for browsing activity in general; active downgrade attacks are mitigated; active, targeted man-in-the-middle attacks are made more observable.

3.4.2.4 Modalities of regulation during transition The shift described here – from occasional security to encryption everywhere – is remarkable in the breadth of re-engineering of technology and re-thinking of norms and practices in a large, diverse and not centrally-controlled group. To give an explanation of that transition might be to explain *why*, what motivated that change in paradigm, what upset the existing handoff and directed the community towards a different one. Identifying

⁷⁶Dramatically: “we had a good thing / you messed it up [...] never again” (Thomson 2014).

⁷⁷In short: UIR, HSTS, the preload list.

the *trigger* (above) is an attempt at such an explanation. Comparing the paradigms themselves and what actors are responsible for security in each is an explanation of *what* the handoff consists in. But another kind of explanation is to describe *how* a change is effected.

In using the handoff model, and as is common in analyses of tech policy, we can refer to different modes of action or different modalities of regulation. For example, from Lessig, we can refer to law, architecture, markets and norms as distinct modalities to regulate behavior, with distinctive properties (1999).

During this transitional time of negotiated re-engineering, the different groups of actors identified use different modalities of regulation; their activities are numerous and diverse. The actors and the modalities they try to use are perhaps not what we would initially assume.⁷⁸

Modalities of regulation interact substantially; there are rarely sharp boundaries. I attempt to group the actions employed during this transition by the modality of regulation that is predominant in each situation. In each case, the action is regulating in the sense that it influences the action of some other actor in our system, separate even from the actions that regulate the ongoing activity within our new or old handoff configurations.

3.4.2.4.1 Market Centralization in the technology field means that many of the companies that compete in one market also play a role in others. Microsoft famously produces and sells operating systems (Windows), and is also a significant browser developer (Internet Explorer) and operates a search engine (Bing), web sites (MSN) and online advertising. That multiplicity means that a browser vendor might use an alternative corporate role to influence a development of the Web. Google announced (Bahajji and Illyes 2014) that sites served over HTTPS would receive a boost in search results rankings.⁷⁹ Given the commercial importance of appearing high on a Google search results page (see: the SEO market), Web site operators had a new incentive to adopt HTTPS, even if it might incur the cost of purchasing certificates or upgrading hardware and software.

Corporate actors weren't the only ones to identify market incentives as important to this engineering change. The Let's Encrypt project was a collaboration

⁷⁸ Whether this assumption is obvious or common I'm not sure, but I think we could typify government actors as using law, corporate actors as using market pressures, engineers as using architecture, advocates as focusing on norms.

⁷⁹ If it weren't so beneficial for end users, we might expect that to fall under anti-trust scrutiny, as when the Department of Justice investigated Microsoft for using its OS monopoly to influence the Web browser market.

between key companies (browser vendors, CDNs) and non-profit advocates (EFF) to establish a new certificate authority (CA), in many ways in direct competition with commercial players. Most significantly, Let's Encrypt provides the certificates necessary for authenticated HTTPS web sites at no cost. Where previously a small web site developer might have had to pay on the order of \$10 a year to purchase and renew a certificate, Let's Encrypt made the process free and mostly automated. This was no doubt an application of direct economic incentives, but it also played a substantial rhetorical role in the larger process of convincing reluctant developers to embrace adoption of a new technology.

3.4.2.4.2 Architecture One debate that illustrates the particular uses of architectural features was the proposal to add “opportunistic encryption” to the HTTP standard. Different proposals might have operationalized that differently, but the suggestion was for servers and browsers to negotiate an unauthenticated encrypted channel even when a certificate wasn't available. The motivation was to provide protection against passive surveillance (this would apply both to the teenager in the cafe and the NSA, in most cases) but without the more substantial guarantees from full HTTPS. In particular, that debate turned on whether Web site operators would consider the opportunistic encryption mode “good enough” and be disincentivized from providing additional security.⁸⁰

Technical standardization proposals have also been used by parties opposed to the spreading of end-to-end encryption. A number of companies provide commercial services that depend upon inspecting and altering communications between Internet users: for example, anti-virus vendors or providers of exfiltration detection and prevention. These “middleboxes” want the capability to intercept these encrypted communications, decrypting them upon receipt, doing inspection for malicious attacks or the departure of sensitive data, and then re-encrypting them. While some end-to-end encryption proponents simply object to this model at all (given the potential for abuse of employees and customers, or alternative methods to achieve those security goals), some vendors have proposed standard ways for explicitly including a proxy as a party to the encryption, breaking end-to-end confidentiality, but maintaining some level of transparency or integrity. As implemented, these architectural means can allow for the continued operation of certain middlebox business models; they also serve a persuasive purpose in trying

⁸⁰Would users be given some UI feedback that the channel was encrypted? If they were, it could more feasibly provide that disincentive for site operators. As some argued, even if users never realized that there was some additional level of security, they could still benefit from it.

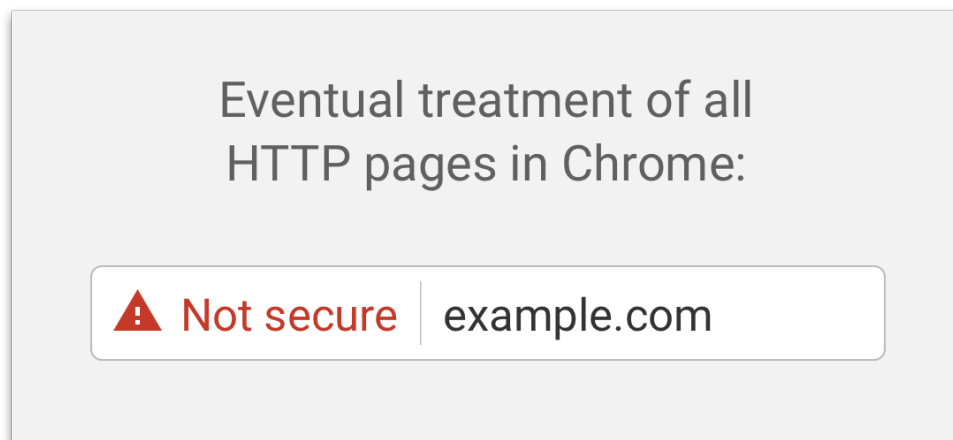


Figure 5: Screenshot of eventual treatment of HTTP in Google Chrome (Schechter 2016).

to promote alternatives that aren't fully end-to-end encrypted, or to provide a negotiating position that end-to-end encryption will be broken in various contexts.

Browser vendors can also use user interface design (which is typically explicitly not standardized across browsers) as an incentive for site operators to adopt security measures. These changes are typically made gradually, but Chrome has also signaled that it will eventually treat Web pages loaded over HTTP as explicitly "Not secure."

That red warning triangle might indicate to users something about the security situation that has long been normal, that there was no technical guarantee. More important for the purposes of this transition, it also provides a visual discriminator that might encourage users who are comparing sites to be cautious or wary of sites that are HTTP only. In that way, the code delivered to the many users of Google Chrome (on in this case, the blog post announcing some future changes in code) can affect market incentives.

3.4.2.4.3 Norms

The IETF community's technical assessment is that PM [pervasive monitoring] is an attack on the privacy of Internet users and organisations. The IETF community has expressed strong agreement that PM is an attack that needs to be mitigated where possible, via the design

of protocols that make PM significantly more expensive or infeasible.
(Farrell and Tschofenig 2014)

These are strong, blunt statements from a technical standard-setting organization. While direct about the values implicated (privacy), the framing is also limited in discussing “technical assessment” and denoting “attack” as a technical term rather than a judgment of malice. Similarly, while this is a community call for addressing surveillance in the design of standards, it is not as strict about specific conclusions as it might be. (There have been discussions of a “no new plaintext” document, but no such strict policy statements have been published.) Recognizing a consensus and describing “strong agreement” among a group is one way to document and encourage a change in norm.

3.4.2.4.4 Laws State actors notably have access to another mode of regulation of action; they can pass laws, rules and regulations and use law enforcement and penalties to encourage compliance. US intelligence agencies have repeatedly called for laws that would more explicitly restrict use of encryption so that wiretapping of Internet communications for law enforcement investigations would be easier. Legislative proposals from the FBI in May of 2013, for example, would have added financial penalties of \$25,000 a day for Internet companies that did not successfully provide wiretap capabilities (Savage 2013).

This is another phase of the “Crypto Wars,”⁸¹ a popular term used to describe debates between law enforcement and Internet companies and civil liberties advocates over the accessibility of encryption to the public. While these are debates over potential legislation, we might also interpret the very public statements of government officials as attempts to influence the norms of design of Internet communications technology.

3.4.3 Using handoffs What do we gain from the handoffs model of analysis for the shift to encrypting Web traffic?

Identifying the handoff in values provides some protection against the naive assumption that a value simply didn’t exist or wasn’t provided prior to its technical implementation. Confidentiality of communications existed prior to TLS or to HTTPS-everywhere, it was just an unsteady assurance, provided by a mix of legal, social and market incentives. Identifying a trigger and a new paradigm provides

⁸¹Or perhaps, as new proposals are about the re-designing of technology altogether, the “Design Wars” (Deirdre K. Mulligan and Doty 2016).

a richer explanation of why this massive re-engineering of a system took place rather than a purely technical one: that a value wasn't present before, and now suddenly was.

In some ways the handoff here is straightforward, and may be a model for security features in many cases: discretion and responsibility is being removed from the end user (or some uncertain assumptions about other participants) and enforced cryptographically. To the end user, this might just appear like simple progress: if only more responsibilities for security vulnerabilities could be taken out of our hands (less constant vigilance about lock icons required, say) and instead guaranteed technically.

But how the handoff is actually accomplished is more complex: it relies on the coordination of many different actors – Web server operators around the world, notably, among others – and a combination of norms, market forces and architectural changes developed the path to the new paradigm. We can look at handoffs as shifting responsibility for a value, but also a triggering event and actions not just within each static paradigm but the modalities that move the socio-technical system between them.

When we apply the same model to Do Not Track,⁸² we'll see a different handoff (not just human vigilance to security guarantee) but also a different set of actions within and between paradigms.

⁸²See Do Not Track, a “handoff”.

3.5 Do Not Track, a “handoff”

3.5.1 An ad that follows you Out to lunch with a friend, the conversation drifts to buying holiday presents. You have been struggling recently to come up with an idea for a gift that will surprise your spouse; your friend recommends a particular brand of watch that you haven’t heard of before. You pull out your smartphone and type the name into the search box; your friend taps the link to the appropriate online store and shows you a couple of the colors he thinks your spouse might like. Lunch arrives, and you put away your phone and put aside shopping plans for now, there are still a few weeks before the holiday.

That evening, you’re sitting on the couch next to your spouse, who mentions a particular news item from the day. Pulling out your laptop, you load an article on the topic and scroll through it; to your shock, you see an ad in the middle of the article for the exact purple watch you were looking at over lunch. Hoping your spouse hasn’t seen it, you quickly click “Back” and open another article instead, and see the same ad. “Oh, were you thinking of getting me one of those?” So much for that little surprise.

For days and days afterward, you keep seeing those ads again and again, on your phone, your laptop, the shared tablet that you keep in the kitchen. All the more frustrating because you’ve chosen not to get that watch after all, once it wasn’t going to be a surprise, but you still see it, in a series of colors, often multiple times in a day. How was it that your phone talked to your laptop, or the watch manufacturer to the different news sites? Who knew you were looking at this particular product and why was that disclosed to your partner? *Who or what was really responsible?*

Could *you* have prevented this scenario? Probably, using existing technology. If you’re aware of this problem and thinking ahead of that possible outcome, you might open a “private browsing” tab on your phone before that first search; when you’re done looking at different watches, your browsing history is erased (along with associated cookies) and that’s probably enough to prevent the “re-targeting” that revealed your shopping plans. Or you could have installed an ad blocker on your web browser at home so that you rarely see ads anyway. Those individual actions may be effective, but is that how we would determine responsibility here? What if the company knew you didn’t want to see those ads everywhere and had refrained from showing them? Or could some part of the system have limited the ads so they only popped up on your phone? Could you tell the advertisers not to customize ads in that way or otherwise control what you see?

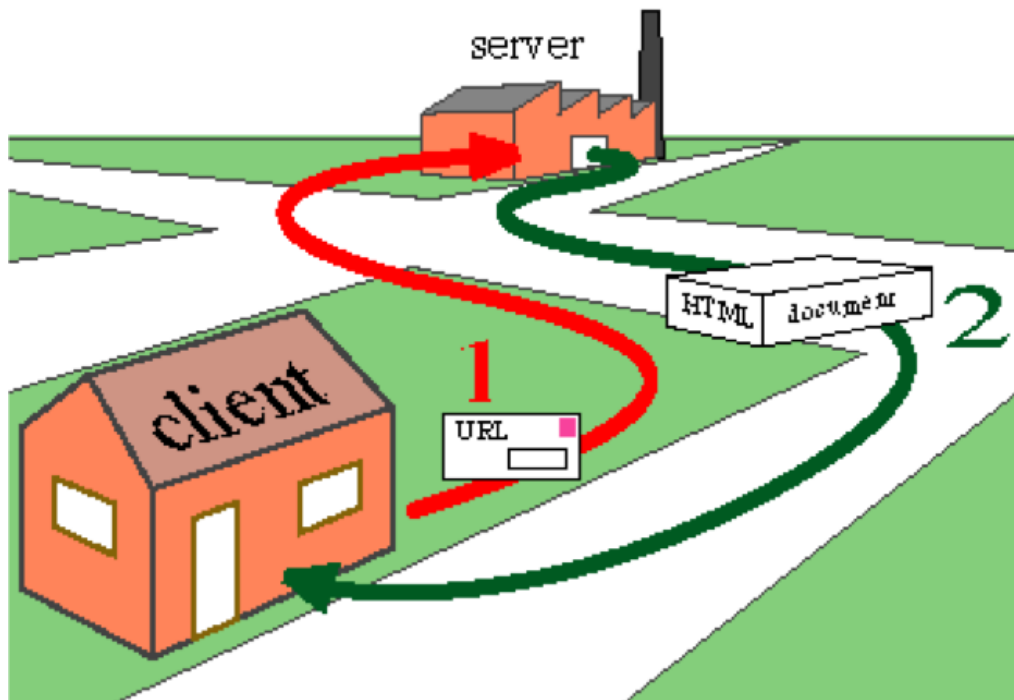


Figure 6: A diagrammatic representation of the Web. Source: CERN.

3.5.2 Handoffs

3.5.2.1 System overview The Web as a **socio-technical system** is complex in both its makeup and function. Billions of end users use web browsers on personal smartphones, laptops, smart televisions, desktop computers at their local library or Internet cafe. Web sites that those users visit are produced by newspapers, governments, corporations, non-profits, individual hobbyists.

In its simplest conception as a user-operated client requesting a Web site from a single server operated by a host, the parties are clearly separable and easily identified. (See the [Web client-server diagram](#).) But in understanding the typical commercial arrangements used for hosting, caching, analytics, market research and advertising, the picture is more complex. (See the [display advertising diagram](#), for one small portion of that detail.)

This more complicated landscape of interactions can also be made somewhat visible in the system of requests for resources that make up a Web page. What I

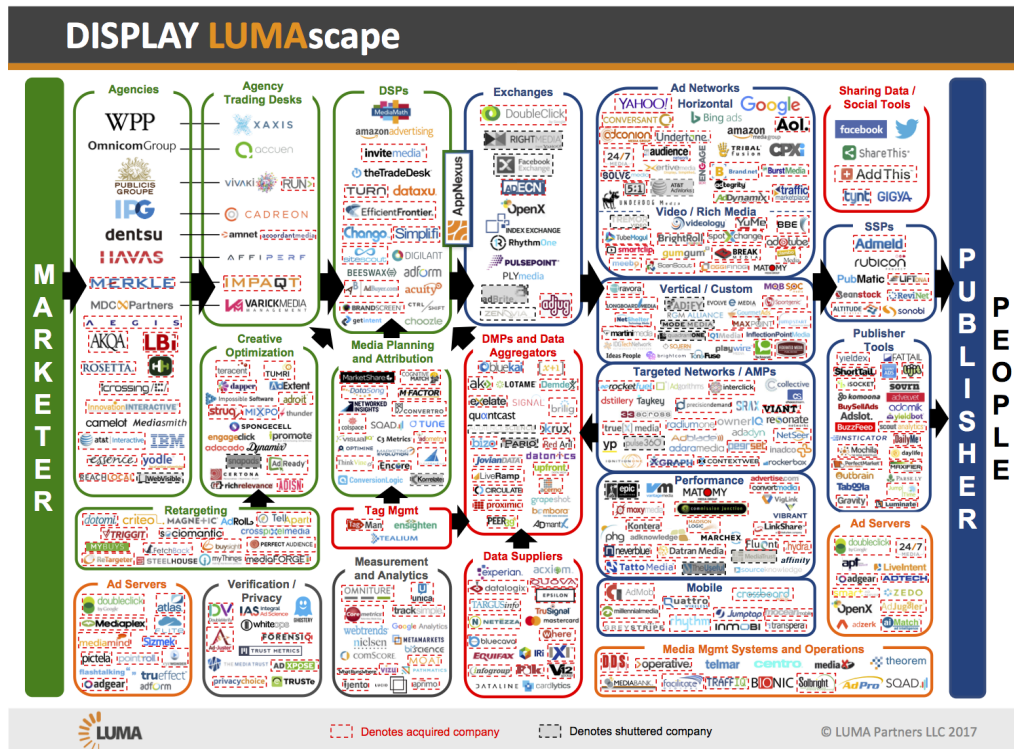


Figure 7: An overview landscape of companies involved in online display advertising; one of a series of popular landscape images from LUMA.

have found to often be a surprise in presenting the technical architecture of the Web to non-technical audiences, your Web browser typically makes a large number of requests to load all the resources that make a modern, graphically-intensive Web page. That same infrastructure is used for many analytics and advertising-related purposes; requests are made, behind the scenes, so to speak, to servers operated by analytics and advertising companies, and those communications include information about the user and about the page the user is visiting.

How we define that complicated interconnected socio-technical system and its scope is itself a challenge. Identifying the active stakeholders may be one guide: open multistakeholder processes typically invite participation (or recruit participation) by groups that are likely to be impacted by changes in a particular design. Engagement in political rhetoric or debate also provides an indication of scope. While participants from ISPs were involved in Do Not Track standardization de-

bates, we saw more involvement and focus on the higher layers of the Internet's design; this was a Web topic. Impact on the public, on a larger and less differentiated group of users, of citizens, is harder to gauge this way; nonetheless, consumer advocacy organizations and political figures (including elected officials as well as administrative agency leadership and staff) became deeply involved in Do Not Track efforts.

The **actors** that make up our socio-technical system then include both technical pieces (Web browsers, networks, servers), the organizational complexity that arrange those operations (browser developers, advertising networks, analytics vendors), legal and regulatory regimes (the Federal Trade Commission, the EU General Data Protection Regulation), as well as people (users of the Web, individuals who participate in technical standard-setting).

3.5.2.2 Handoffs between actors Collaborators have defined a **handoff** as the transfer of a function or the responsibility for some value from one actor to another between two different configurations of a system (Deirdre K. Mulligan and Nissenbaum 2020). Exploring that shift in responsibility can provide some insight into the political and societal consequences that are too often considered unforeseen or uncontrollable.

Within every configuration of a socio-technical system, there are distributions of responsibility and functionality – sometimes explicit, but mostly implicit and often misunderstood – among different actors. It can be tempting to think of security in network communications as a value provided purely by technical measures (encryption, say); however, deeper analysis would typically show that security is provided in part by technical measures and in part by legal enforcement, organizational practices, and community norms. In trying to locate responsibility for privacy in our ad re-targeting example, we will come across those rough edges between different actors in the current system, and how the proposed and actual re-configurations of the socio-technical system change how the responsibility for that value is distributed. Understanding why those transfers occur is useful in providing a full explanation of how technological changes affect society.

The history of Do Not Track is so fascinating because we see an attempt to make the distribution of responsibility between technical and legal regulation explicit and because we see an attempt by activists to embody a value in a technical design while explicitly not enforcing that value technologically. These potential handoffs stand in stark contrast to the more unidimensional shifts seen in the high-level trends of automation or privacy-by-design. And seeing this as a handoff

better captures the complexity beyond simple comparisons between technical and legal regulation.

3.5.2.3 Diversity of goals As the functionality of this sociotechnical system depends on the complicated interactions of many different actors, the goals that are implicated for the Web as a sociotechnical system also vary.

Many people use the Web for personal communications: checking their email, posting messages to social networking sites, reading and writing blog posts. Commerce is a common set of uses that is especially relevant to this example: companies provide services for sale; people buy both digital and physical products; online advertising is widespread; media companies provide entertainment services. As we might see more specifically looking at other illustrative examples, there might be very different goals in mind for parties like intelligence agencies or state actors, that may be orthogonal to or in opposition to the goals of many individual users of the system.

We could also identify goals from the stated purposes of designers of the system and its components. The first Web page ([Berners-Lee 1992](#)) sets out a succinct and exciting goal for the project:

The WorldWideWeb (W₃) is a wide-area hypermedia information retrieval initiative aiming to give universal access to a large universe of documents.

Universal access to a large universe of documents gets at the goals of the originators of sharing information, about ongoing scientific projects but also other topics, that can be easily searched and browsed, and implying both retrieval but also easy writing and publication. Berners-Lee even uses the language of the system “aiming to” accomplish that singular goal. However influential that original stated purpose might have been, or might still be among people intimately involved in technical decision-making regarding the Web, it’s clear that this system is now complex in a fundamentally different way, that no single person or small group of people has control over the function or the direction. The multistakeholder model of technical standard-setting – through which new functionality for the Web is debated and agreed on – reflects the variety of independent but connected stakeholders that are affected by and jointly implement the Web.⁸³

⁸³For more, see [Internet Standard-Setting and Multistakeholder Governance](#).

That the socio-technical system does not have a singular, agreed upon goal is useful in understanding the tensions in how to distribute responsibility for a particular function or what values (and what particular interpretation of those values) should be designed for in different configurations.

3.5.2.4 Paradigmatic changes How can we determine responsibility for providing privacy while browsing the Web, as in our initial motivating example? To illustrate the different distributions of how privacy protection is provided within a system, I describe three different system configurations representing three paradigmatic approaches: first, a cumbersome self-regulatory opt-out regime combined with a set of browser cookie controls; second, a proposed co-operative approach with expressed and respected preferences; and third, an active arms-race of ad and content blocking.

3.5.2.4.1 Traditional notice and choice Privacy concerns related to the profiling behind online behavioral advertising have been present as long as that business model has been widespread. In the US, the Federal Trade Commission helped negotiate privacy practices with industry self-regulatory bodies, as part of its initial series of reports and actions on online privacy in the 1990s ([Federal Trade Commission 1998](#); [“Self-Regulation and Privacy Online, FTC Report to Congress” 1999](#)). The notice and choice model was implemented, in part, through “opt-out cookies” – using the same basic technology (HTTP cookies) typically used for tracking user activity, an interested user could visit a page in their browser that would set opt-out cookies for each of a potentially large number of online behavioral advertising profilers and that cookie would be sent on subsequent interactions. Promises were made by participating online advertising companies to comply with those self-regulatory codes, including to limit the display of behaviorally-targeted ads. These opt-out cookies have been criticized as cumbersome and ineffective ([Dixon 2007](#); [Leon et al. 2012](#)): the process of clearing cookies (which you might do for privacy reasons) would effectively opt the user back into profiling and behavioral advertising; cookies might be set to expire and the participating companies would change over time, so users would need to regularly re-visit and re-install opt-out cookies; and cookies were specific to a single browser, so the same process would need to be applied repeatedly across browsers and across devices; finally, the scope of the privacy choice was unclear or unsatisfying, you might still have your browsing information collected by the same parties using cookies and just not see the targeted advertising until the opt-out cookie expired.

Browsers typically provided a user interface for viewing and clearing cookies, and some experimented with plugins to provide transparency into the different interactions with online services that could track user behavior. But determining which cookies were required for functionality (for account logins and commenting interfaces and shopping carts) and which might be for tracking browsing activity across sites was typically infeasible for the user. User education efforts suggested clearing cookies on some regular basis, but doing so also implied the inconvenience of logging out of sites. Third parties developed browser plugins for blocking trackers, or for blocking the display of advertising altogether. Techniques began to be developed for “re-spawning” cookies; taking advantage of browser bugs, browser plugins or configuration details to maintain identifiers of a user even when cookies were cleared.

In this paradigm, user privacy (at least for the re-targeting example in the anecdote above) is available to the user through cumbersome or uncertain actions on their part, with the legal and normative backing of industry trade associations and a regulatory body, or potentially through technical means, although those means were already being outmaneuvered.

3.5.2.4.2 DNT While we might typically identify activists in the area of online privacy as focused on technical solutions, Do Not Track was proposed as a solution that used technology but did not rely on technological enforcement. Rather than continuing an arms race of cookie-management/browser-fingerprinting, an extremely simple machine-readable signal was to be standardized. Browsers and devices could communicate that signal to other parties online (including both the web sites you visit and the additional parties involved in online advertising and other services), who could comply with the user’s expressed preference not to be tracked. Adoption by online parties is voluntary, or at least not enforced by the technical protocols themselves.

In this proposed paradigm, privacy is available as a simple choice to the end user, and that choice is expressed through their browser software and enacted through a similar mix of self-regulatory industry action and the potential for regulatory enforcement. DNT’s technical mechanisms are designed specifically to allow for enforcement of a user preference through a combination of consumer regulation, industry self-regulation and software changes. How those choices are enacted, and whether the user understands whether their expressed preference is respected is not technologically enforced, but left up to that combination of private organizational ordering, legal mechanisms and technical designs.

3.5.2.4.3 A new arms race Currently, DNT standardization has been completed without widespread adoption by online services and major online advertisers have indicated that they will not modify tracking behavior in response to a user's expressed preference. Industry trade associations and self-regulatory groups have not further developed any alternative browser-based tools. Up to this point, browser vendors have maintained a Do Not Track setting for users, but have also developed more nuanced technical tools for blocking requests or cookies. The use of ad blockers has increased, in add-ons, modes and dedicated browsers. Some publishers rely on vendors to detect ad or tracking blocking and impede or block access to their published content.

While the focus of this analysis has been over distribution of responsibility for the value of privacy, motivated by privacy concerns regarding collection of browser history and disclosure in alternative contexts, this phase of ad-blocking arms race notably involves other values. Ad and tracking blocking software is designed for and advertised as promoting a broader range of values – performance improvements, better security or a less distracting reading experience – in addition to, or instead of, the preservation of privacy.

In this paradigm, competing software design changes – on the client-side and the server-side – impact user privacy, but also security, performance, access to content, and web site business models, with changing implications that are hard for users to measure but may be more visible.

3.5.2.5 Modes of action In modeling handoffs between configurations, we consider not only the modalities of regulation – markets, law, architecture and norms – used by the various actors within our socio-technical system but also other properties of their actions – whether they are visible or invisible, expressive or coercive – which are described as the **mode** of action.

Of particular relevance here is that we can distinguish between the actions within each of the three paradigms as well as actions used to negotiate or move between those paradigms.

For each paradigm, what are the prominent actors and modes of action and how do they interact?⁸⁴

3.5.2.5.1 Modes of action within traditional notice and choice Most prominently featured in the traditional notice and choice paradigm (see [Traditional no-](#)

⁸⁴We could also organize these by the modality of regulation – markets, law, architecture and norms – as I've done in the Encrypting the Web handoff discussion.

tice and choice, above) are the self-regulatory arrangements: negotiations between the FTC and NAI and certifications and audits of online behavioral advertising organizations. These negotiations are typically private, don't involve direct consumer representation and may be unknown or invisible to the end user.

This opt-out paradigm relies on certain technical arrangements as well. HTTP cookies are re-used for organization-by-organization opt-out communications, and a Web application both explains the opt-out process and allows for setting those opt-out cookies. These are architectural measures that are implemented and controlled by participating online advertising companies, using the existing technology of cookies as it's implemented by Web browsers; the cookies are expressive signals (implementations typically didn't delete other cookies the advertising networks may have set) but the signal is both set and received by the same party. Opt-out cookies are explained and configured through a web page operated by self-regulatory industry groups, rather than a browser setting or control.

The arms race over this tracking activity, especially in leading up to Do Not Track discussions, features different presentations of controls to users by different parties. Browsers provided cookie clearing as a user-initiated method for inhibiting tracking and educational efforts (a kind of norm-setting) suggested clearing cookies as a part of digital hygiene. Optional add-ons for blocking tracking or blocking ads saw some small levels of adoption. Cookie clearing and management sees a technical response in techniques for correlating activity without relying on the persistence of HTTP cookies, including browser fingerprinting and cookie respawning. While user controls have a direct effect (deleting records stored on their local devices), the arms race makes the effects increasingly obscure and uncertain.

Many technical measures are not self-enforcing mechanisms. Some tools provide increased transparency (including the Lightbeam plugin, pictured) about tracking connections between sites, or the numbers of trackers present. That's an architectural modality of regulation, but it works primarily to persuade or influence other actors, whether it's end users, businesses or regulators.

3.5.2.5.2 Modes of action for DNT proposals Do Not Track combines some of the properties of opt-out cookies and direct blocking tools. A DNT header is expressive rather than coercive or self-enforcing: it merely communicates to some other party that a user prefers not to be tracked. But it's also a communication mediated in a different way than a trade-association-managed opt-out cookie: users have the option to select DNT in their choice of browser

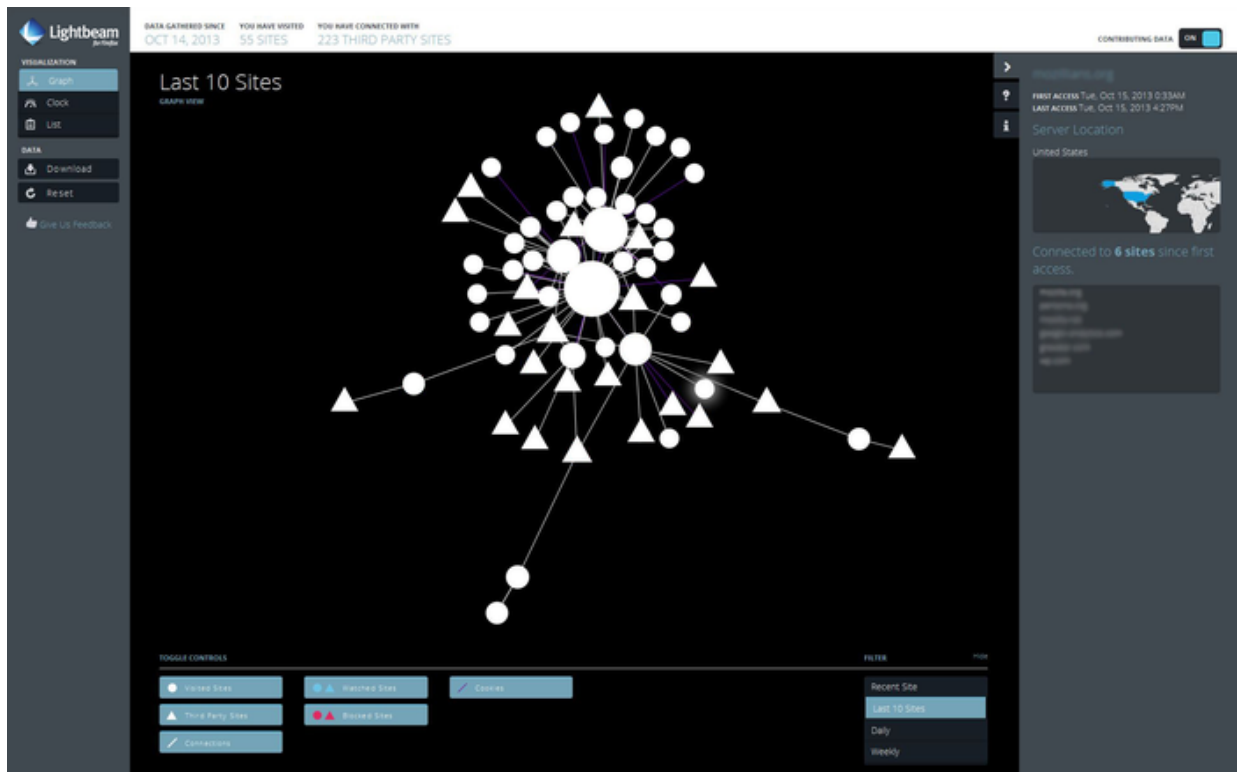


Figure 8: The [Lightbeam](#) (previously “Collusion”) plugin visualizes common third-party connections from visiting multiple sites.

software.

DNT as proposed relied on negotiations, if not formal agreements. The standardization process attempts to find consensus among the different parties that might use the DNT header about its meaning and how to comply with it. The W3C standard-setting process is open to a larger and wider variety of stakeholders and its discussions are publicly archived, but this is still largely invisible to the end user.

Enforcement of DNT could happen through distinct means: legal requirements may require or incentivize complying with user preferences in some jurisdictions; statements of compliance may be enforced through trade regulations (for example, FTC enforcement); self-regulatory groups could provide industry agreements and trade associations or other groups could provide external audits of those commitments. Some proposed tying blocking measures to assertions of DNT

compliance: tools that block cookies or other tracking mechanisms could refrain from those blocking measures for parties that respond to an expressed preference. That may be a real-time negotiation on behalf of the user (“I’ll let you collect some data, so long as you promise to respect my preference not to combine data about me on different sites”), but mediated through expressive signals sent by an online service and client-side measures to block cookies.

3.5.2.5.3 Modes of action in blocking and counter-blocking As an implicit or explicit response to the delays in standardization or the lack of server-side adoption of Do Not Track, browser developers have integrated more sophisticated technical responses to tracking. Attempts have been made to systematically block or limit storage while minimizing breakage of popular embedded functionality.⁸⁵ Machine learning and other heuristics are increasingly used, beyond the simpler and more static allow and deny lists that were previously proposed. Heuristic, learned and list-based approaches are less direct in the sense that a user-facing control has more complex implications, but the semantic description and the likely implications may at the same time be more comprehensible. “Block tracking scripts” both implies something more complex but also more accessible than deleting a cookie from a particular origin.

As publishers (especially news organizations) increasingly employ paywalls – limits to the number or selection of articles that are available before a user is prompted or required to subscribe – there has also been an increase in blocking access to content for users who are detected as blocking online advertising or tracking. This uses both technical and market measures: the blocking can be accomplished technically, but using pay subscription as an alternative provides a financial incentive to allow advertising and tracking. Market incentives also apply to the browser vendors: performance and privacy protection can be selling points in the competition for users, while the possibility of sites blocking access with a particular browser could cause users to switch.

While the technical means in the DNT paradigm are expressive, blocking and counter-blocking attempt primarily to be self-enforcing or directly effective. The visibility and transparency of these actions is also different: blocking technology can be obscure or opaque (in much the way that tracking technology long has been); paywall prompts are a more explicit, expressive message, and issued from

⁸⁵For example: Firefox’s Tracking Protection and Safari’s Intelligent Tracking Prevention and Storage Access API.



Figure 9: A billboard in New York City advertises the Firefox browser based on outwitting online tracking, November 2016.

the publisher to the user rather than from the third-party ad networks. The effects of this [shift in responsibility](#) are discussed further below.

3.5.2.5.4 Actions that influence the movement between paradigms The previous sections identify the actors and properties of their actions within each of three possible paradigms of our socio-technical system. But those configurations don't exist in parallel or come into being deterministically. We can also observe the actions taken to influence the handoff between different configurations of a system, involving many of the same actors and a diversity of modalities of regulation and modes of action.

One prominent potential starting point in the timeline for Do Not Track is a report from the Federal Trade Commission staff recommending development of a standardized Do Not Track mechanism. This is a notable instance of a government

agency actor not using law or rules as its modality of regulation, but rather using communication as a form of norm-setting. Throughout the DNT process, FTC has used diplomacy and encouragement of stakeholder participation, rather than rule-making or bringing enforcement actions.⁸⁶

Participants describe Do Not Track debates as an especially political process, both inside and outside “the room.” Lobbying and other kinds of influencing are about setting or changing norms through direct or directed communications. That can involve closed-door lobbying of government officials, certainly, but also public messaging, aimed at users, at companies in the industries involved, or at administrative or legislative representatives. Participants cite references to emails/videos regarding interpretations of a chair’s comment at a particular TPWG meeting and a campaign to tie targeted advertising to saving kidnapped children.

Technical and architectural measures are used as means of influencing discussions. Consider two software *patches*⁸⁷ introduced during particular moments in DNT standardization debates: a proposed change to Firefox’s cookie-setting policy to accept cookies only from visited sites; and a proposed change to Apache’s default configuration file to ignore DNT headers sent by Microsoft Internet Explorer. Ultimately, neither of these patches was accepted by the corresponding open source software project, but the demonstration of the technical approach was an attempt to influence market forces. While these may not be unique in the history of software development, persuasive software patches are certainly idiosyncratic.⁸⁸ This form of communication is also limited in its accessibility: it requires programming expertise, technical reputation or both to contribute these changes, and indeed it takes some technical expertise and understanding of open source software development methods to understand (or translate) the implications of such changes.

⁸⁶The FTC’s choice of regulatory actions depends in part on statutory restrictions, historical limitations of administrative rule-making and an approach of engagement, topics covered in great detail by other scholars (Hoofnagle 2016; Bamberger and Mulligan 2015).

⁸⁷A patch is a self-contained proposed change to a piece of software code and is the typical method for introducing, discussing and adopting new changes to collaboratively developed software. The name comes from the older practice of patching over punch cards or paper tape to change a piece of software that was already distributed.

⁸⁸Another example might be the development of plans for 3-d printed firearms: while some might try to develop and use such weapons, it’s commonly accepted that their promotion is an attempt to discourage gun control regulation (Manjoo 2013).

3.5.3 Using handoffs What do we gain from the handoffs model of analysis for the different Do Not Track configurations? In identifying the complex set of actors at different scales; their choice and the mode of their actions; and, the variety of shifts in responsibility that are considered, we can see what is distinctive about Do Not Track and the debate over user privacy of Web browsing activity.

3.5.3.1 A network of actors and actions Analyzing the socio-technical system as a network of actors and their use of different modalities of regulation can uncover the potentially complicated tensions between various forces at play. This kind of analysis is more familiar in tech policy and science and technology studies as in Actor-Network Theory (Latour 2007) and code-is-law (Lessig 1999). This is just a first step in describing handoffs, but identifying the actors and modal properties of their actions – hard or soft, expressive or self-enforcing, transparent or opaque – can make the implications more explicit for analysis.

An in-depth understanding of Web architecture shows not just the endpoints (the abstract client and server) but also parties that are, abstractly, in the middle, or lower-down: the Internet Service Provider used for connectivity by both the user and the online service; middlebox vendors providing services within enterprises or on in-home networks; the different companies involved in developing and maintaining the user's device, operating system, Web browser and DNS resolution; the parties involved in delivering the diversity of Web pages and their embedded services, analytics, advertising, behavioral tracking and content delivery. Companies are not easily separable into those categories, most notably because many large technology companies compete in multiple areas: Apple sells hardware as well as developing operating systems and a Web browser; Google has the most popular Web browser but the vast majority of its revenue comes from online advertising. Even within the category of online advertising there is diversity of positions: there are different sizes of online advertising networks and different services that different companies provide, and those ad networks and ad technology vendors are distinct from the advertiser itself, that is, a company that has paid in order to show a text or graphical ad for their product or service.⁸⁹ That complex network of organizations makes it harder to identify the “sides” in a debate – or even who

⁸⁹It's interesting in this DNT and online privacy context that people who refer to “advertisers” often mean those who sell advertising, like Google and its AdSense network, and not organizations that buy ads, like Coca Cola or car companies, say. Consider the difference between Clear Channel, which might own the large billboard down the street, and Nike, whose ad featuring Colin Kaepernick you might have seen on that billboard. Increasingly, tech companies like Apple and Netflix, are also prominent buyers of outdoor advertising like those billboards.

can speak for or adequately represent what group – or create a simple mapping of who wants what or where a compromise might be. Browser vendors and online publishers might seem like natural mediating parties: browsers might have a closer connection to users and publishers typically have legal agreements and technical measures in place with embedded third parties providing advertising, analytics and other services, but the level of visibility and control that each has is unclear, and our paradigms haven't previously put responsibility on those companies.

It can be tempting to identify categories of technology with the large companies that sell or operate those systems, but in fact there are individual humans who develop software while employed by Google and individual humans who attend meetings with the FTC or visit congressional offices. There may be studies where identifying the individual backgrounds and experiences does not add significantly to an economic analysis of the market positions of the employing organizations, but this is not such an area. Particularly in the Internet field, individuals move between companies and take their experiences and positions with them. Individuals also have multiple roles beyond just their primary professional employment, including their roles in open source software projects and in technical standard-setting bodies. In DNT discussions, roles within companies (engineering vs. sales or product, say) mark a distinct grouping separate from and sometimes orthogonal to employing organization or industry.

This example demonstrates not just a diversity of actors, but the somewhat unusual actions (which vary in their modality of regulation and other modal differences) from our cast of players. In our timeline, the Federal Trade Commission is prominently cited, but not for taking an enforcement action or proposing rules, but recommending a technical mechanism and encouraging standards development. Consumer advocates engage not so much in political lobbying, but join in the technical standard-setting process and provide technical expertise and proposals. Members of Congress send a letter of comments to the World Wide Web Consortium on a public mailing list. Microsoft, a developer of operating systems, a popular Web browser and engaged in online advertising and online publications, makes a prominent default setting proposal. Advertising trade associations are perhaps more conventional in engaging in political lobbying, but perhaps novel terrain in public relations criticisms of non-profits or Web browser businesses.

3.5.3.2 Shifts in responsibility Specific to handoffs, describing the movement and distribution of responsibility can better explain the impact of decisions and changes that might otherwise be seen as value-free. In this case, we are considering

how responsibility for privacy over how data about a user's browsing is collected, shared and disclosed and how that responsibility might be redistributed. The movement between the different paradigms might not be confused for value-free, given the controversy or impact of the different configurations. But the shifts of capability and responsibility are significant and perhaps distinctive in the arena of tech policy. The traditional notice and choice paradigm leaves responsibility unallocated: neither technical guarantees nor regulated arrangements provide a particular sense of confidence about a value like privacy. Instead, as noted in the opening vignette, the end user could execute control⁹⁰ if they implemented a set of uncommon technical changes or abstained from using the Web altogether. One response to such a situation of identified inadequate privacy or security protection is to move the discretionary capability away from (or take the burden off of) the end user and instead to provide a technical assurance: for example, a technical system that blocked all data collection that could be used for profiling and behaviorally-targeted advertising. Another response is to set a norm (perhaps bolstered by law, rules or self-regulatory arrangements) for some backstage actors to provide enough of an assurance to the user that they don't need to be concerned with a technical arrangement that they don't understand or can't control: for example, laws, rules and self-regulation could prohibit retention of user browsing data or its use for targeting advertising.

Our story here differs from these typical paths. Advocacy and regulatory actors called for a technical mechanism, but not for technical mechanisms that provide guarantees, automatic enforcement and a human-free assurance. Instead, DNT is a technical mechanism for communication of user preferences, rather than traditional notice about business activities, between the user and a subset of other parties. This maintains the opt-out metaphor preferred by businesses and some US policymakers, but with some fundamental differences. Browsers present the choice and information about it to the end user, and can do so in a variety of ways, and users have a new method for communicating with those embedded and often invisible third parties. This is a handoff, but not one that simply removes both capability and assumed responsibility from the end user: instead, it increases

⁹⁰This example does not speak to the "notice" part of "notice and choice." I don't know that any user has any such capability to understand the technical means behind how ads are tracked and displayed; I've never seen a user successfully use self-regulatory notice icons for that purpose, for example; meanwhile, rumors about how behavioral tracking works and are basically incontrovertible, as anyone knows who has tried to explain to their friends that smartphone microphones aren't constantly listening to their in-person conversations in order to later target an ad for display on Instagram.

communication and makes a kind of shared sense of responsibility between users, browsers (also known as user agents) and the plethora of analytics, advertising and tracking partners.

The new blocking arms race is perhaps more analogous to the security/encryption case. There is still a new handoff, a shift in responsibility: browsers are taking a more direct role in blocking trackers, ads or other resources. These new approaches are less mechanical and less user-directed than the less-widespread alternatives discussed for previous paradigms: there's mostly not a direct list, or a choice of parties to block or unblock, and settings are more likely to be automatic or tied to some other mode rather than user-initiated. Instead, browser developers provide data and algorithms for ongoing identification of tracking and blocking in ways that aren't anticipated to interfere with user-desired functionality. The resulting arms race situation does increase the visibility of the situation for the user, in the case of paywall notices described within the main content of a Web page, and requests for users to provide data explicitly, or become paying subscribers, or to change their browser mode or preferences. Whether and how this situation benefits or diminishes privacy depends on how we conceive of that value. Users of these blocking tools might have less data collected about them but there's little predictability about what tracking is happening when as the different parties try to work around each other's tools. Explicit negotiation with sites over privacy and payment was one of the intended outcomes of Do Not Track as an opt-out mechanism: it makes those tradeoffs more apparent to the user, but might also contribute to different parties collecting different user data (like billing details).

3.5.4 Distinctiveness in handoffs In our initial example, we saw responsibility for privacy as amorphous and uncertainly placed: who's responsible for this ad that follows you and what can be done about it? By considering different paradigms and the diverse, distinctive actions within and between them, we can evaluate different handoffs of that responsibility between a complex network of actors. Each paradigm – notice and choice, Do Not Track, blocking and counter-blocking – has distinct implications for the value of privacy: whether users have control or rely on others and whether those controls are accessible, effective and enforced technically or through some combination of policies.

The handoff model also helps us analyze the particular properties of the actors and actions within and between those configurations. Debates over DNT included software patches that were effectively persuasive rather than architectural. And Do Not Track is distinctive in being a proposal for a technical mechanism to support

user privacy that is expressive rather than self-enforcing and a system that relies on broad multi-party cooperation.

4 A Mixed-Methods Study of Internet Standard-Setting

This chapter describes the complexity of where, how and by whom Internet standard-setting happens and how I seek to study it, as an active participant, using semi-structured qualitative interviews and analysis of communications data. My goal is for this explanation to be useful both in interpreting the results of this research and in contributing to the research field of Internet governance.

For both participation and research, Internet standard-setting can at once be both surprisingly open and frustratingly opaque.

- participation is encouraged by: open process, extensive archives, ethos of individual participation
- participation is inhibited by: expertise, costs (time, money), history with a group

[Setting the scene](#) provides context for this openness and opacity with an example of the standard-setting work mode and its networked nature. The openness of participation and nonetheless the substantial barriers to it make for significant questions of the legitimacy of technical standard-setting as a model of governance or regulation.⁹¹ For the researcher, access to the standard-setting community reflects the tension of those participation trends. We have a rich corpus of participants, conversation and design, done in a relatively open and well-documented way. But there is also a maze of bureaucratic detail, technical jargon, pre-established personal relationships and outside-the-room activity, spread across many groups and hundreds of organizations in different sectors and geographic areas.

[Studying standard-setting at different scales](#) considers those properties of the technical consensus standard-setting process and the combination of methods for understanding the community and the implications for these multistakeholder processes in addressing debates over values such as privacy. Each of those scales is detailed in the following sections. To address that challenge, I have used my personal involvement and participation to inform my inquiries, interviewed a sample of participants privately to gain an understanding of their diverse perspectives and used automated analysis of mailing list archives to identify and measure broader

⁹¹See the earlier chapter [Internet Standard-Setting and Multistakeholderism](#) and [Doty and Mulligan \(2013\)](#).

trends across a larger community. Finally, I present the [ethical framing](#) of studying up and the protections for research subjects.

4.1 Setting the scene

I recall the first standards meeting I attended, showing up in no official capacity and with no particular affiliation.⁹² Ten or fifteen people sit around a U-shaped table in a nondescript hotel conference room in Santa Clara, California; they are mostly white men, engineers at different tech companies. Someone at the front of the room is the chair, and projects a list of items onto a screen that the rest of the group faces. Other chairs are available, not at the table, around the edges of the room, where people less committed to this particular meeting can sit in, watch, maybe participate. I take one of these seats; at one of the breaks maybe I introduce myself to a couple of people. Everyone in the room has their laptop open. While there is a discussion happening in person, there is also an online chatroom – known to all the participants, to some remote attendees who couldn’t make it in person, to some people who have an interest in this group but are currently in other conference rooms in the same hotel – with active discussion. At times, there is relative quiet in the room while everyone is typing and reading what others are typing. (Yes, this experience feels bizarre at first, but you get used to it.) Whenever someone is talking in the room, someone else (the currently-designated “scribe”) is taking notes on their statement, attributed to the speaker, in the IRC room.⁹³

The conversation is loosely organized around a list of issues (that projection onto the screen in the room) related to the document being discussed:⁹⁴ the chair or editor asks a question about how something should be phrased or explained, and people in the room provide brief opinions; some back-and-forth, some through an organized person-by-person queue. The meeting is small and fairly casual; as the youngest, least experienced person in the room, I still feel comfortable chiming in on occasional points. Some of the questions are answered right here – after a brief back-and-forth, it seems like everyone is in agreement, the chair points that out and summarizes, the issue is “closed” and the resolution is recorded in the

⁹²A vignette of a typical scene is provided to help the reader grasp the basic structure of this process; this is not based on field notes.

⁹³The [minutes of this particular meeting are publicly available](#), recording most every statement, including my own questions.

⁹⁴In this particular case, it was the draft specification of the W3C Geolocation API, being developed from 2008-2009 or so, that will enable web sites to ask to determine the user’s precise latitude and longitude. Many W3C and IETF meetings follow a similar agenda pattern.

chatroom by the scribe. Someone else in the chatroom adds some notes with more detail, or types out who is responsible for implementing the change. But in many other cases, the question can't be fully resolved right here and now: someone needs to investigate a technical detail further, or an argument between two participants needs more fleshing out, and so it's noted that the conversation will be taken "to the list."

"the list" is the group's mailing list; a group like this relies on this piece of automated email-based infrastructure. This is by formal organizational policy in most cases: IETF and W3C create hosted mailing lists whenever they charter a new Working Group (or indeed many of the less formal groups as well). Meetings cannot be easily organized without such a broad, accessible communication channel, and a group without out-of-band asynchronous electronic communication can't do the discussion that makes up the work of a group made up of people living in different countries, working at different companies and participating either intensively or just occasionally. The mailing list's address is available on a corresponding web page about the group and widely advertised for feedback on any standards documents. Typically, the mailing list is public: anyone can subscribe, anyone who can convince the system they aren't a spammer can send a new message that will be distributed to the full group, anyone (subscriber or not) can read through a Web-hosted archive of every message ever sent to the list in the past. Messages range from very short ("+1") to thousands of words long. There are, to anyone not familiar with this kind of work, a lot of them. Long threads with very detailed arguments about any issue considered by the group, automated or bureaucratic messages describing technical changes or distributing the minutes of past meetings, casual personal conversation or complaints about some piece of software or another that are met with a reminder about the intended scope of the mailing list. List conversation can be friendly and informal or, at least as often, brusque and insulting. Prior to attending that conference room "face-to-face," I'd sent a few messages to the group's list; some would receive responses from an interested party or someone trying to gather support behind a particular idea, others would be ignored.

But even looking at these two "sites," distributed and wide-ranging as they are, would be too blindered to understand all the points of conversation between the formal participants of a standard-setting process, much less to capture: the debates within organizations; the business and policy discussions between firms inside and outside the same field; the relationship of companies and regulators in different jurisdictions; or the effects of software out in the world and how it's used. I recall describing this research project on privacy in standard-setting to

an important policymaker (a non-participant stakeholder, in the terms of this research) who asked me, pointedly, was I going to limit my investigation to the standard-setting groups themselves? (Someone else in the room quietly shook their head “no,” urging me to avoid my obvious blunder.) Trying to play the good academic, I responded with something about focusing research on a limited scope for the purposes of finishing a dissertation; this was quietly received as a sign of my apparent obliviousness.

In laying out the illustrative cases in this work, I have followed discussions that formally take place in standard-setting fora, but also describe interviews with participants and non-participant stakeholders and cite relevant press writing and other announcements broader than just the working groups.

4.1.1 The networked site While not traditional in the sense of historical anthropology, this is also not an entirely novel environment. For example, Coleman has written about the free and open source software movement by researching at week-long conferences and reading extensive mailing list archives (Coleman 2012); Some Internet ethnography has tried to focus on the properties of the virtual site as a place (relying on metaphors of “cyberspace,” for example), the communities that exist in that place and how to conduct research “in the field” in those settings (see lists of citations in Davies 2012). But more relevant to the distributed, wide-ranging and on-line/off-line communication styles of standard-setting is Burrell’s argument of a networked site with multiple entry points and connections (2009), Kelty’s view of “distributed phenomena” (2008) and Hine’s ethnography of mediated interactions and the “richness and complexity of the Internet” (2000).

Rather than marking a bright and arbitrary border around the site, I have tried to follow the participation in the distributed Internet standard-setting process *where* it happens – ranging from formal in-person meetings of groups with a specific membership, as well as teleconferences and mailing lists – as well as *who* is involved – from formal leadership to those non-participant stakeholders who observe or influence without being in the room. At the same time, to focus my inquiry, I have tried to focus on some core settings and then expand out to more peripheral involvement to supplement that study. Following the cases described in *Privacy and Security: Values for the Internet* above, I have set as a core group, the Tracking Protection Working Group who debated Do Not Track standards at W3C. While that group does have, in some ways, a formal membership list, meaningful participation also expands out to people who joined teleconferences, in-person meetings or mailing list conversations, those non-participant stakeholders who




Figure 10: Sketch of core to periphery of stakeholder groups around Do Not Track, showing “wedges” of participation.

followed the process or influenced it in some ways, and further out the casual observers or even just the affected parties who had no awareness of it. By necessity, different stakeholders and sectoral groups will be larger outside the group than they are internally, even where represented internally. In the DNT work, some groups have deeper core participation (for example, advertising companies and consumer advocates) while others (for example, policymakers) are more peripheral. These wedges of depth of participation extend outwards: another way to picture participation.

The debate over encrypting Web traffic doesn't have as singular a core standards body locale. The same debate was ongoing in specific working groups, security area groups, in IETF plenaries and within and between companies involved in implementations and deployments of software at different layers of the stack. Even so, we can see conversations happening in Internet and Web standard-setting at the center of a set of concentric circles that encompass firms, advocates, policymakers and users. There is overlap in the participants in standardization around Do Not Track and standardization around Web encryption even though the technical details of those projects are quite distinct.

This diagramming is not intended to value the importance of one group or one setting over another; as biased and interested in the impact of standards as I am, I still wouldn't call it the most important part of the development of the Internet and the Web. Indeed, research subjects are explicit in denying the framing of standard-setting as the most essential step and public writing from standards experts also emphasizes that point, as described in [Internet Standard-Setting and Multistakeholder Governance](#). Instead, these diagrammatic maps of the standard-setting process show how this study of technical standard-setting processes is situated and how these distributed working groups are connected to others.

4.2 Studying standard-setting at different scales

In studying Internet technical standard-setting, even scoped to Internet and Web consensus standard-setting around the values of privacy and security, I am faced with the challenge of grappling with this diverse, distributed, networked site.

In order both to validate findings from intensive qualitative investigations and to identify the character and causes of apparent trends present in quantitative data, this study takes a mixed methods approach to investigation. This site involves, significantly, individuals with their lived experiences, dynamic interactions between people, and organizational structures that connect large industries. To encompass those scales, this research also includes different methods to gather insight at different scales of study. Different research projects can contribute by including just one of these methods, but my central argument focuses on the interaction between individual participation and larger organizational settings and so this work attempts to examine those different components.

Table 1: Applying different methods to different subjects and scales of study

Methods	Subject of study	Scale
self-reflexive	lived experiences	micro
qualitative	interpersonal dynamics	micro/meso
quantitative	organizational, high-level trends	meso/macro

How can these different methods at different scales inform one another? Here are three examples from my ongoing research:

identifying areas for closer investigation Quantitative analysis of a large dataset like the full corpus of IETF mailing list archives allows measurement of social network properties like closeness centrality to identify important figures that are highly connected to different subgroups.⁹⁵ Identifying central people or groups in a large community can help identify key people to interview or establish what organizational roles are worthy of more intensive observation and inquiry.

validating and measuring disparities My personal experience at standards meetings prompted the question of demographic imbalances in standard-setting group participation, observing male-dominated in-person discussions or apparent reticence of participants from East Asian countries. Quantitative estimates of gender in participation on mailing lists provides a point of comparison (do computer-mediated communications have the same effects?), a way to validate (does the disparity apply across multiple contexts?) and to identify potential variables that might affect the distribution (do some work areas show less demographic imbalance?).

explaining the effects of interventions Explaining interventions is an especially challenging task for research. Quantitative analysis can provide comparability, but suffers from confounding factors or wide variations in interpretation. Qualitative research can provide rich description, but shies away from causal explanations or broad external validity. By using both quantitative analysis of the trends across hundreds of documents and a qualitative understanding

⁹⁵See this research notebook on IETF participation for the methods, data and initial conclusions from calculating closeness centrality: <https://github.com/npdoty/bigbang/blob/ietf-participation/ietf-participation/IETF%20Participants.ipynb>

of reading documents and talking to authors we can better explain and contextualize the effect of mandates and guidance on the presence and significance of security considerations sections in standards (Doty 2015a).⁹⁶

Subsequent sections of this chapter describe the specific methods and tools used to gather and analyze data from different sources, roughly grouped in the same order of self-reflexive, qualitative and quantitative methods described here. Integrating those methods as I suggest⁹⁷ is an ongoing challenge for me, but one where I hope to make a contribution.

4.3 Researcher position

As a participant as well as a researcher, I aim to use my own perspective, including the challenges of working in a diverse multistakeholder setting. I have approached privacy in standard-setting as an active participant myself — not a putative detached observer — and that will inevitably be apparent in my work. Research requires reflexivity (Watt 2007) about the research process, my own subjective experience and the effects of my involvement. While “reflexivity” is used in many different ways, what I intend here is methodological reflexivity (M. Lynch 2000) — awareness of my own beliefs and experiences (because they are unmediated) and awareness of my multiple roles within the groups of study themselves.

Not unlike many subjects of this research, I have held many different roles (at different times and simultaneously) with respect to development of the Internet and World Wide Web. I have been a user of the Internet and the Web since I was first given a one-hour tutorial at my local public library in small-town Virginia (circa 1992) and managed to discover a Star Trek fan page. Like many in my generation, my technical training was mostly self-taught; in middle school (the mid-90s) I started a web design “business” with a classmate (we never had paying clients, although he later founded an online community called Reddit), and learned enough HTML and JavaScript to make images change on mouseover. While computer science classes were accessible in high school and college, they never covered Web technology; I taught myself PHP and explored the privacy violations present by running `tcpdump` on misconfigured switched networks. After a short

⁹⁶While this has been a part of my research practice on standard-setting, as in the cited paper, this dissertation work does not report on measurements of the effects of interventions.

⁹⁷This model of mutually-informing scales of study was also presented at the Protocol to the People event at the Turing Institute (Doty 2018).

stint in software engineering at Microsoft, I first encountered Web standards as a Berkeley graduate student, following mailing lists and writing workshop position papers and research reports regarding geolocation privacy. From 2011 through 2015, I was employed as part-time staff at the World Wide Web Consortium (with “Privacy,” rather than a job title, on my business cards), managing the Tracking Protection Working Group (TPWG) and the Privacy Interest Group (PING), work that ranged from being a job for one to two days each week up to entirely consuming all my waking thoughts. As staff, I recruited participants, handled process and logistics, organized meetings and events, provided technical support and responded to press inquiries, but also participated in the discussion, debate and design work of the groups themselves. In 2013, I started editing a guidance document for mitigating browser fingerprinting (2019) and in 2014, I took on the role of Editor for Tracking Compliance and Scope in DNT (2019), roles that I continued as an unpaid volunteer after 2015.

Throughout these different roles, I also identified myself as an academic, a researcher affiliated with UC Berkeley with an interest in privacy and technical standards, and continued (where I could find the time) to publish research in journals and workshops. Having those multiple, overlapping roles (user, amateur Web developer, W3C Staff, academic, Team Contact, Editor) complicates my experiences and my position. But such role diversity is also not exceptional among standard-setting participants themselves, who switch employers, job titles, rhetorical positions and stakeholder groups while maintaining a connection to a standard-setting body.

Having a stake in the outcome might seem like a violation of the neutrality expected from:

- a) a researcher,
- b) someone writing a standard for compliance, or
- c) the staff of a standard-setting body.

In each case, I believe that purported detached neutrality is neither plausible nor constructive to the aims of the commons.

A common point of confusion in explaining the standard-setting process to press was the apparent contradiction between standards both regulating the behavior of, and being debated and developed by, the implementers of those standards. “Isn’t it like the fox guarding the henhouse?” This question arose most often when noting that initial editors of DNT specifications included representatives of Google and Adobe, but the question and confusion applies to the process of de-

veloping specifications more generally, especially because editors do not typically have ultimate decision-making authority. Consensus standards can function not despite but *only because* they are developed by the impacted groups that are the implementers.⁹⁸ For practices to be voluntarily adopted and practically informed, it is constructive to have stakeholders as authors and collaborators. Scholars have argued that standard-setting is possible only because of the shared motivation towards a common goal (Cargill 1989).

Standard-setting bodies vary in the roles of staff and the different kinds of neutrality that they practice. W3C itself has a mission, the apt, if anodyne: “to lead the World Wide Web to its full potential by developing protocols and guidelines that ensure the long-term growth of the Web.” Staff are expected to share that goal and work towards that mission, which is explicitly non-neutral. But there is an expectation that the Consortium is neutral with regard to the members, many of which are direct market competitors. Staff have very different backgrounds and participate to different degrees in the groups they facilitate, but have considerable discretion within the broad mission and can take vocal positions in standardization discussions. This is to some extent inevitable and to some extent a particularity or historical tradition of Internet standard-setting: that individuals have perspectives and express them in a way that is distinct from organizational priorities.⁹⁹

Finally, as a researcher, deep involvement with a particular perspective seems contrary to the traditional position of a detached observer as in the model of the pith-helmet-wearing anthropologist documenting primitive tribes.¹⁰⁰ There is ongoing qualitative research on technical standard-setting communities using participant observation and interviewing from researchers observing these groups; I look forward to seeing their results. However, detachment can sometimes assume an impartiality that does not exist, where reflexivity allows us to recognize our positions as researchers (our personal perspectives, our effect on discussions, how our methods respond to community dynamics). In terms of access, rapport

⁹⁸See “The Consensus Standard-Setting Model” in *Internet Standard-Setting and Multistakeholder Governance*. That some impacted groups are necessarily well-represented, however, does not mean that all impacted groups are, which is a key question for the legitimacy of techno-policy standards (Doty and Mulligan 2013).

⁹⁹For more on the distinctive role of the individual in Internet standard-setting, see: “The Internet and Requests for Comment” in the chapter *Internet Standard-Setting and Multistakeholder Governance*, “A network of actors and actions” in the case *Do Not Track*, a “handoff” and *Individuals vs. organizations* in the findings.

¹⁰⁰For some audiences in the area of qualitative research, the absence of a neutral detached position is so widely accepted as to need no explanation, while for others this assumption may be important to rule out.

and nuanced understanding, I believe there are distinct advantages that my deep personal involvement can bring. Mine is necessarily a unique perspective, and one I hope can bring the reader more deeply into a complex setting.

My own subjective experience can be an advantage in connecting with subjects in the community and understanding their lived experience (Sandelowski 1986, discussing the advantages of subjectivity for confirmability). My first-hand familiarity with the challenges of multistakeholder participation has been useful in establishing rapport during interviews and in identifying (via experience) and confirming (via interviews and conversations) emotional responses present in the community. However, because my own positions are clear to stakeholders, that could inhibit candor where a participant chooses not to expose disagreements. In addition, where stakeholders have identified me as antagonistic, that may limit access altogether. As a researcher working on privacy, my position that privacy is an important value to be supported in the design of the Internet and the Web has been well-known, and while that view might be common in the abstract, there are surely other participants who identify me as too focused on the value of privacy, or not focused enough, compared to other values.

I stay aware of the limits of relying on my own experiences, lest the study become confessional.¹⁰¹ Personal experience can seem all too vivid, but lacks the replicability or investigability of rigorous, systematic research. In studying these multistakeholder groups, which almost by definition include people with very different backgrounds and perspectives, using my own experiences as canonical would be a mistake. I have tried to use personal experience in the form of vignettes or self-reflection to illustrate a perspective and to illuminate settings that may be unfamiliar, but rely on qualitative and quantitative methods to analyze the community and processes. Noting the issue of reflexivity does not automatically deflect all concerns, but this awareness should exist throughout my presentation of this research and in evaluation of my methodology.

4.4 Interviewing

Even where documentation is extensive and participation can be directly observed, understanding the internal views of members of a community can be difficult.

¹⁰¹“Here is just one example of the total wrongness of something I tend to be automatically sure of: everything in my own immediate experience supports my deep belief that I am the absolute center of the universe; the realest, most vivid and important person in existence.” — Wallace (2005)

To gather insight into that emic perspective, I have conducted semi-structured interviews with various participants (and some non-participant stakeholders) in Internet standard-setting.

These interviews were guided to gain insight into both the people and their personal perspectives and the standard-setting process and how they perceived it.¹⁰² I began interviews with questions about participants' backgrounds and their roles in their organization; I inquired into personal views on privacy – how they define it, what kinds of privacy concerns they identify; and then I asked them about their experience with technical standard-setting, when privacy came up in those conversations and how they saw their role. For those involved in the Do Not Track process in particular, I asked about their particular goals for that process, how they perceived debates that took place and how they viewed other participants.

I conducted 27 interviews over the course of this project, not evenly spaced between late 2012 and late 2019. That distribution was largely driven by this researcher's varying time that could be dedicated to data collection as my direct participation and employment took up less space. As a result, the interviews don't attempt to show a comparative assessment at a single snapshot in time, but include both ongoing and retrospective viewpoints.

4.4.1 Dimensions for sampling Understanding the perspectives of a diverse community, or a community that is at the intersection of various groups rather than a single cohesive or homogenous setting, provides challenges for the sampling process. While we cannot *a priori* know all the variety among our potential research subjects, we can sample in a way informed by theoretical considerations. Among those: multistakeholder groups specifically prompt the question of how different stakeholder groups are represented and operate in such a process.

Can we neatly divide the participants of, say, the Tracking Protection Working Group into a clear faceted classification of distinct stakeholder groups? No, but we can nonetheless identify important apparent distinctions. Mapping different and overlapping stakeholder groups is feasible based on my working experiences with Web standardization and analysis of membership lists. We might also profitably use mailing list participation as a proxy to confirm or expand the representation of different groups, to the extent that we can evaluate affiliation. Research subjects can confirm or reject that sampling frame. While I have access and understanding of

¹⁰²The full interview guide I used for privacy in the standards process is included in the [appendix](#).

different participant groups to sample from, I also ask participants to recommend particular people to talk to as a form of “snowball sampling.” This is done less for convenience and more to get the participants’ own views of what groups or perspectives might be missing; in order to maintain the privacy of participants from other participants in the study, names are asked for without revealing who has already been interviewed.

For W₃C standardization as a whole, I have tried to represent in this diagram the different overlapping groups of stakeholders, their intersections of what they represent and where they align and their levels of participation and influence in W₃C processes. This working sketch is based on a review of W₃C organizational membership and my personal experience with W₃C organization and working groups; as such it is only one personal view out of potentially many different ones, but my hope is that it could give outsiders some idea of what sectors are present.

In addition to sampling different stakeholder groups, demographic differences in participants is important for our ethnographic study to explore the effects that standard-setting process might have on demographic representation or the participation of different subgroups. There are many demographic dimensions that may be relevant to questions of legitimacy over the design of Internet protocols. Because tech communities face prominent controversies over sexual harassment and discrimination in employment contexts, gender is one such area of interest. Only 15% of my interview subjects were women: similar to the gender balance in tech firms and in Internet standard-setting groups, but still very far from proportional.

Because this is a theoretically-informed sample, we might also choose to oversample certain groups of potential importance. As leadership was a theme identified early (from interviewees not in leadership positions), this research has tried to particularly include those with formal or informal leadership roles.

While studying a particular standard-setting working group may allow for a precisely-defined scope of membership, this study also seeks to include non-participant stakeholders and peripheral participants. Peripheral participants – sometimes involved and sometimes not, or attending meetings but not vocal – might provide an outside perspective and insight on why people choose not to participate.

Finally, sampling can be explicitly used to mitigate biases, either in the researcher or in potential access. One such danger is that as a non-confrontational person myself it might be especially easy to speak with a significantly skewed sample of people who are generally agreeable or who share perspectives or goals with me. Statistical representativeness is not a goal of this method, but not sampling at all from entire subgroups with a particular, distinct perspective would harm the

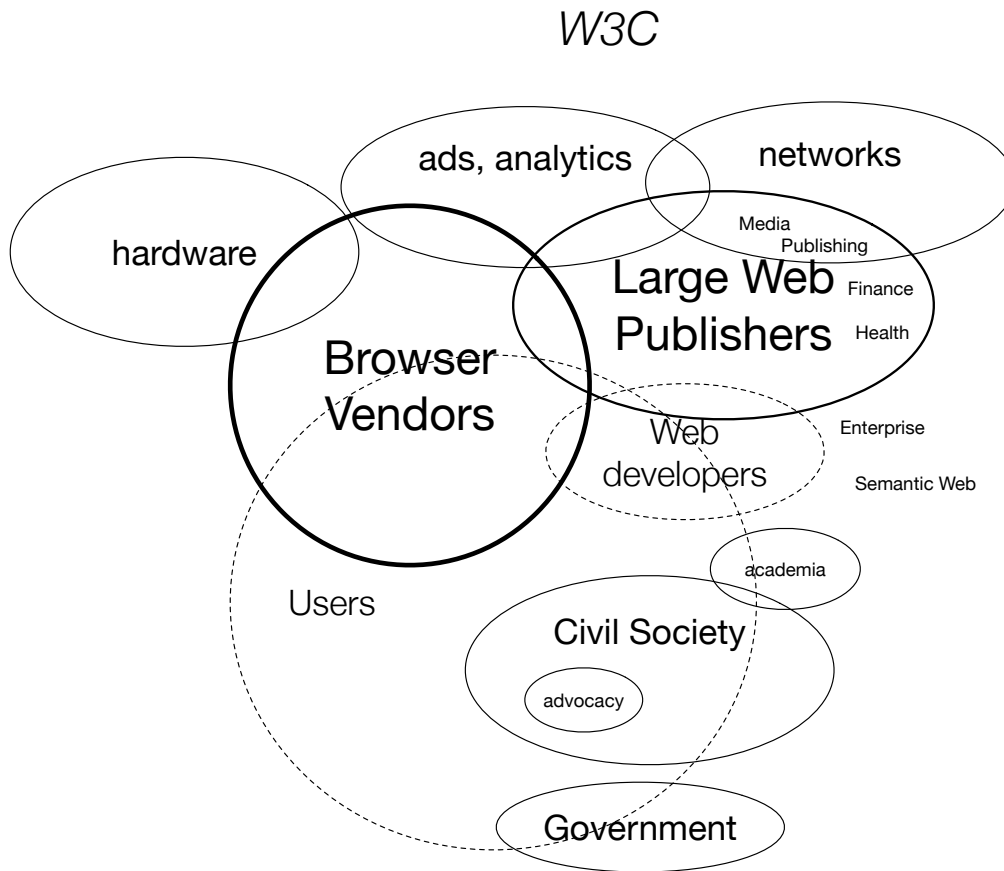


Figure 11: Mapping of overlapping stakeholder groups at W3C.

breadth of understanding.

There is a subtle but important distinction between being more-or-less friendly to work with and being more-or-less supportive; indeed, these might be orthogonal dimensions. For my own purposes in identifying this variety in stakeholders and confirming the intentional diversity of my sampling, I developed this two-by-two matrix of Tracking Protection Working Group participants based on my own experience. For this purpose, I believe using my own experience to be especially apt, as I'm attempting to counter any internal preference for similar or agreeable subjects.

For the dimensional axes: the vertical axis shows the spread from being “friendly” (easy to work with, e.g.) to “difficult” (more likely to have unpleas-

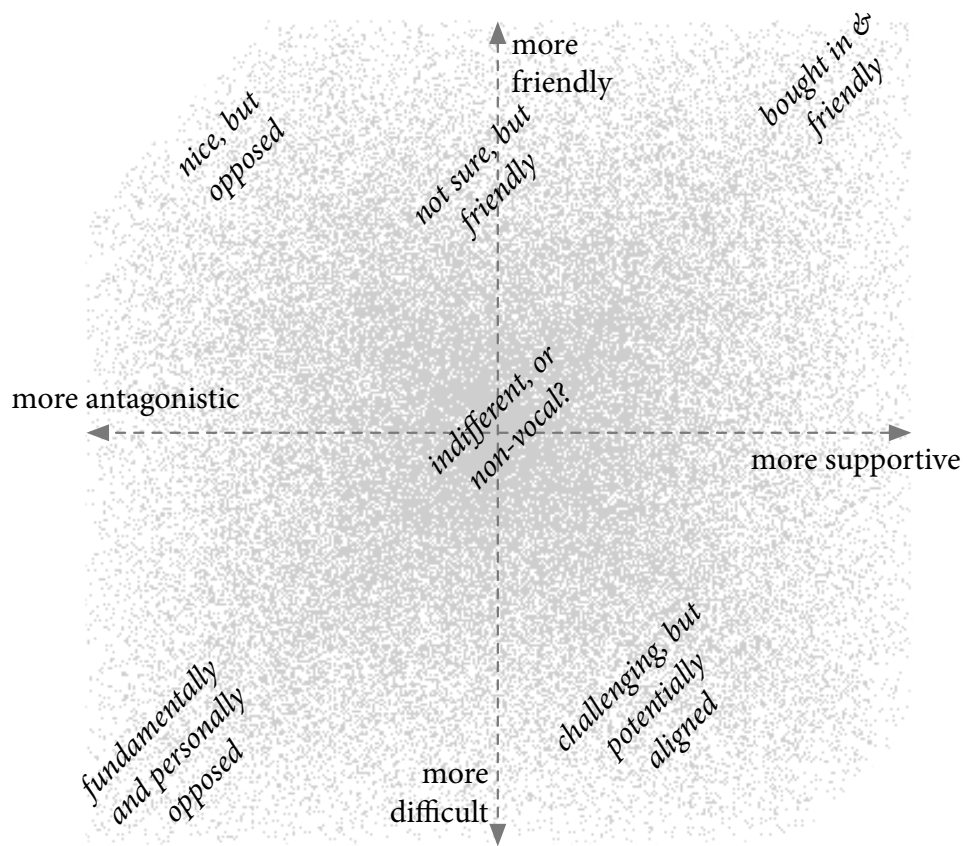


Figure 12: Sketch of distribution of orthogonal dimensions of difficulty-antagonism, with personal notes for the DNT standardization process.

ant interactions); the horizontal axis is between “supportive” (shared the basic goals of the process) and “antagonistic” (opposed the existence of the process or the stated goals). There are certainly people in the friendly/supportive quadrant: people bought in to the process and happy to collaborate. And the direct opposite quadrant is also easy to identify: people who were opposed to the Do Not Track process in every way and who seemed to personally dislike me as well. This group may be difficult to access. That the dimensions are orthogonal (or at least substantially distinct) is supported by the presence of people in the other quadrants. A substantial group of people were regularly professional and nice, always happy to talk about professional and personal topics but were nonetheless substantially opposed to Do Not Track or to W3C processes. And while there may not be the most extreme examples, there were also people known to be difficult to work with who were nonetheless significantly supportive of, or potentially aligned with, the standard-setting process. As reported in the Findings, interviewees confirmed these dimensions as orthogonal and even posited similar additional ones about participating in good or bad faith.

Beyond those dimensions of diversity, I’ve attempted to collect data up to a point of meaning saturation (Hennink, Kaiser, and Marconi 2017).

Regarding the saturation parameters, this work has weights on both sides of the scale, as highlighted in the diagram. The heterogeneous population; emerging, conceptual codes; and theoretical interest all suggest a larger sample. But having access to thick data (and other sources of data, although this mixed and multi-scale method isn’t considered) and iterative sampling allow a relatively smaller sample for saturation.

4.4.2 Coding, memoing and writing through Interviews with these participants, who have been largely engaged and candid with their experiences and expertise, provides an incredibly rich corpus, reflected in the transcripts that pile up in encrypted disk images. So much is covered, and there are so many threads to pull at, that it can be overwhelming.

I’ve followed a practice of initial coding and focused coding (Strauss and Corbin 1990; as cited by Lofland et al. 2006) to identify common themes and important contrasts from that larger corpus.

In addition to coding to capture the topics and terminology raised by participants, I want to capture also the holistic perspectives that I hear from subjects. To do this, I write brief (approximately 1-page) memos after completing coding an interview, sometimes supplemented by the handwritten notes I took during

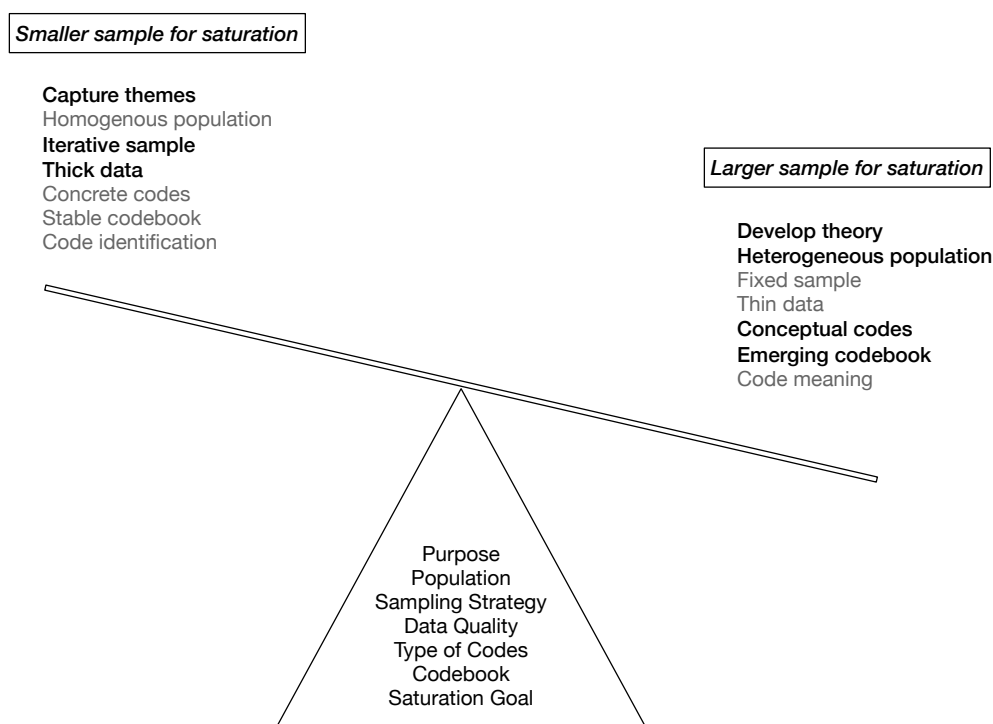


Figure 13: A modification of the [Hennink, Kaiser, and Marconi \(2017\)](#) parameters of saturation diagram, highlighting the particular factors used in my study.

the interview, to describe the ideas that stood out to me – a particular motivating argument or novel idea, say – as I reviewed the interview experience.

Without trying to replicate every point that the research subject is making, memos provide a brief but more open-ended mechanism to collect key points or insights that might not be easily represented as a single term or phrase for a code. At this slightly higher level, I can write down perspectives that might not be obvious in the interviewee’s words but still come through in my reading of the conversation. Theoretical memos capture “momentary ideation” ([Glaser 1978](#); as cited by [Lofland et al. 2006](#)) of my own thinking on reviewing an interview during the coding process.

To draw important insights from the corpus of interviews and their associated codes (hundreds after the multiple rounds of coding), I’ve identified common themes that appear in codes across multiple interviews and clusters of related codes into key themes to drive deeper analysis and my write-up of findings. In

reviewing the quotes from different interviewees that touch on that same theme, I can get a sense of the variation and pull out quotes that are illustrative of a typical viewpoint (expressed several times by different people) or a distinctive viewpoint (one that sharply contrasts with other views).

As this methodology is qualitative, I don't have the statistical backing to show numerically that one view is most common or that some arguments are significantly more widely held than others. Instead, my goal is to show the existence of important themes among the participants and non-participant stakeholders and then to use the subjects' own words to help describe that experience for the reader.

4.5 Analyzing communications

Internet standard-setting is rich with communications artifacts: mailing lists, meeting notes, drafts, revisions and published documents. In this work I have focused on group mailing list archives for communications data that show group discussion and interactions, but there is a real opportunity for our research field to investigate these traces as a supplementary method more generally.¹⁰³

4.5.1 Mailing list analysis More than any other single "place," Internet standard-setting has happened on mailing lists. While not contained in a single geographic location, the mailing list functions in a place-like way, along the lines of what we see in critical geography's analysis of place as opposed to space – as Massey argues, made up of social interrelations and flows of people and communication (1994).

There are other places of interest – important decisions are made in face-to-face meetings or persuasive conversations that happen in private settings at an office or bar – but the majority of argument, debate, discussion, positioning, presentation and reasoning in these groups takes place in email fora.¹⁰⁴ These lists are places in the sense of containing and mediating these interactions, even though the participants are geographically and temporally dispersed. That these mailing lists are typically publicly, permanently archived makes them rich sources of retrospective study and analysis.

¹⁰³See, for example, the concept of trace ethnography (Geiger and Ribes 2011).

¹⁰⁴This is a contingent, historical conclusion rather than a necessary or normative one. There are some groups that use chatrooms for a larger fraction of discussion and recent use of GitHub and software-development lifecycle tools for issue tracking are becoming more common in Web standards work.

4.5.2 BigBang

BigBang is a toolkit for studying communications data from collaborative projects.¹⁰⁵

Collaborators at UC Berkeley, Article 19 and the University of Amsterdam have developed BigBang as a collection of tools for analyzing traces from open source software and Internet governance groups; this collaboration is supported by DATACTIVE, a research group focused on data and politics. Independent researchers pursuing their own projects have collaborated on this tool because of commonalities in the communications tools used by these software development, standards development or decision-making groups – all typically use archived mailing lists to develop community and discuss their work. Identifying who participates, how they participate and how the structure of these groups affects their work is valuable to our group of social science researchers even though – or perhaps especially because – different collaborative communities use these online communication tools in distinct ways.

BigBang has been used for collection and analysis of mailing list archives and Git version control repository information. Functionality has been developed for the following forms of analysis (with examples of specific measures or data considered):

- traffic analysis (messages over time)
- demographic analysis (gender, affiliation, country of origin, etc. of participants)
- social network analysis (centrality, connectedness, assortativity)
- content analysis (trends in word usage)

Those different forms of analysis allow for responsiveness to different classes of research questions. My hope is that we can learn from practice what kinds of data is appropriate to what kinds of research questions and what the practical challenges and feasible solutions are in studying mailing list data. That kind of work can set researchers up for a wider range of future projects along these lines.

Traffic analysis can illustrate patterns of activity: that might include trends across standard-setting fora altogether or the typical pattern of a working group. That activity can also show community responses to exogenous events, as I've

¹⁰⁵<https://github.com/dataactive/bigbang>

explored with privacy and security activity after especially relevant Snowden disclosures about Internet surveillance (Doty 2015a).

Demographic analysis can provide evidence on who is participating, relative levels of participation between different subgroups and how demographic variation changes over time. These analyses are especially useful in answering the prominent questions about access to technical standard-setting fora and conversation, which are important to establishing concepts of fairness and legitimacy when these standards have values implications. I use a combination of email sender metadata and manual annotation to estimate relative fractions of gender within and across mailing lists and I believe similar techniques can (and are, and will) be used to evaluate disparities in participation by sector of affiliation and region of origin.

Network analysis capabilities allow for testing of hypotheses of network formation, for example. Benthall determined that open source development communities did not show the kind of preferential attachment model – a “rich get richer” form of social network development – that has been observed in several other formations of links, as in between websites (2015). Where mailing lists provide important settings for group communication and social network development, we can use this functionality to measure and compare macro-level properties of these groups. Network analysis can be useful in identifying individuals who play leadership or connecting roles around particular topics or between particular groups.

Content analysis differs from these other types in actually looking “inside the envelope” at the text contents of email messages.¹⁰⁶ Measuring how often words arise can help us see trends in where and how concepts like privacy and security are being discussed. And connecting content analysis with network, demographic or traffic analysis can provide evidence of who is bringing up particular values, how concepts migrate across different groups and when topics see more or less attention.

BigBang developers use a mailing list for discussion, and git and GitHub for sharing source code and coordinating work; it has gone through debates on intellectual property and licensing very familiar to open source software and Internet standard-setting. Use of the tools under study and borrowing a working model similar to the communities under study seems fitting, a scientific analog of “recursive publics” (Kelty 2008).

¹⁰⁶While I have explored some use of content analysis for detailing how terminology is getting used or spreading across mailing lists over time, this dissertation does not report on that work or make use of content analysis.

4.6 Ethics

This work examines people (participants in technical standard-setting, non-participant stakeholders), processes (multistakeholder fora; government, civil society and corporate decision-making; software engineering practice) and architecture (the Internet and World Wide Web). Ethical considerations guide how I have conducted research at each of these scales including in how I direct my inquiry, in protecting human subjects and in handling publicly accessible data.

4.6.1 Studying up The literature of anthropology has in the past called for more “studying up” — investigating those groups in society that are rich or politically powerful and not, as had been a trend, focusing study on communities that are vulnerable, historically marginalized or foreign in the sense of being from non-Western cultures (Nader 1972). The motivations for this shift are both ethical and scientific; ethical in the sense of not overly objectifying and limiting the scope of inquiry of problems to those most vulnerable; scientific in the sense of not missing an entire part of society as an object of study.

The ethical impulse to study the culture of people who wield power, and specifically that power held by technical expertise and exercised through the design of influential and immovable technical artifacts, is an essential motivation for this work. The architecture of the Internet and the Web, like many software constructs, have implications for fundamental human values (Nissenbaum 1998); those seemingly technical design decisions are inherently political (Winner 1980), and, like the highway overpasses of Long Island that cast in concrete the impossibility for public buses to reach recreational beaches, have simultaneously long-lasting but hard-to-see impacts (Caro 1975); the processes used for protocol design decisions have opportunities for governance, but also serious open questions for procedural and substantive legitimacy (Doty and Mulligan 2013).

In the study of science and technology, studying up may mean studying the designers: software engineers, developers, user interface designers and “makers” of all kinds, as opposed to studying the larger mass of users whom we typically identify as having less control over these powerful technical decisions. And designers of Internet protocols and Web standards may also constitute an “elite” (Marcus 1983): with agency, some exclusivity and power as their decisions and market status will often influence the decisions of other engineers and technology companies who build on the Internet and the Web.

4.6.2 Interviewing human subjects To the extent that this research studies individual participants and stakeholders, it qualifies as human subjects research. In particular, semi-structured interviews conducted with standard-setting participants and non-participant stakeholders are aimed not only at understanding the implications of the standard-setting process but at learning about the perspectives, backgrounds and motivations of those individuals.

These conversations are kept confidential to encourage candor from the interviewees and to limit any personal or professional harm that could come from disclosing details of those individuals' perspectives or participation. The names and organizational affiliations of participants is typically not disclosed in this work; quotations are provided with context about the type of participant or organization, but do not include details that could be used to directly identify an individual or their affiliation. Some participants were willing to have quotes directly attributed to them and provided specific consent for that point; their quotes are attributed where I conclude that it provides useful context to the reader.

A research protocol for conducting these semi-structured interviews was reviewed by UC Berkeley's Committee for the Protection of Human Subjects (CPHS, the local Institutional Review Board) and was considered exempt from further review¹⁰⁷ because of the minimal risk of harm in conducting confidential interviews with a non-vulnerable population.¹⁰⁸ After updating with a different funding source and project title, the protocol underwent a much more intensive review process, which resulted in longer (but not more informative) consent forms, updated practices for encrypting data at rest (which were applied to all previous interviews as well), and longer data retention requirements (based on interpretation of federal funding guidelines which I believe to be mistaken). That protocol was approved after "Expedited" review.¹⁰⁹

4.6.3 Mailing lists Mailing list analysis also includes collecting and analyzing the communications of human subjects. For this project, list analysis is used both to study the participants and to study the processes and tools of these groups. Because these mailing list archives are collected and publicly presented for the purpose of review, and participants are typically directly informed of this before

¹⁰⁷For UC Berkeley's CPHS, "Exempt" does not mean that no review took place or that no review is necessary, but that a research protocol as described contains minimal risk and so does not need to be reviewed by a full IRB committee or yearly updates/reviews.

¹⁰⁸CPHS Protocol #2011-11-3796.

¹⁰⁹CPHS Protocol #2018-03-10819.

sending a message to such public lists, Berkeley's IRB provided guidance that no human subjects research review is necessary for this collection and analysis. This is not reducible to an argument that no publicly-accessible data can have any ethical implications for research (an argument which I do not support); these archives are made publicly available specifically for the purpose of access and review by others, including non-participants, and that status is typically well-understood by participants.¹¹⁰

Out of politeness, mailing list crawlers were configured to access list archives at no more than 1 request per second. Local analysis of list archives requires making local copies of all the messages of those public archives; those copies may be shared with other researchers in machine-readable form. While those copies are typically identical to the publicly-available archives, we are not currently making those archive copies themselves publicly available. In some rare instances, the hosts of mailing list archives will remove messages from the archives even after they have been publicly distributed.

¹¹⁰W3C mailing lists, for example, send an automated reply to any new sender to a publicly-archived mailing list explaining that status and require an explicit form approval by the user before a message is distributed to the list.

5 Findings

5.1 Connecting research questions to themes

The initial chapters of this dissertation set out both a theoretical lens and a pair of high-level research questions. First, what are the impacts of multistakeholder techno-policy standards-setting processes on resolving public policy disputes for the Internet? And second, how do the designers of Internet protocols view privacy and how do their views ultimately affect the privacy of Internet users?

The semi-structured interviews I conducted with participants and the quantitative social network analysis I have explored about technical standard-setting bodies provide a plethora of themes of interest; it has been fascinating for me, but not every interesting theme or finding can be included here. Instead, I have clustered themes together and elaborated on those most emphatic in the findings that provide insight for the research questions. There may not be a sharp boundary between those results that speak to the multistakeholder process and its impact on privacy and those that speak to the ethical impacts of engineering and participants' conceptions of privacy, so each section includes internal references to draw those connections.

[How standard-setting process accommodates, succeeds and fails](#) explores themes related to the multistakeholder, rough consensus process of technical standard-setting itself. I lay out different stages of a standard-setting process, each of which can be an area of success or failure depending on one's goals for the process. The argument in Chapter 1 described the potential for boundary organizations to accommodate diverse perspectives in a productive way and this section shows the challenge and importance of accommodating good and bad faith behavior, antagonism and disputes among heterogeneous participants.

[Anti-trust and competition in Do Not Track and setting standards for online privacy](#) elaborates on a particular incident identified in several interviews and supplemented by my own experiences and documents from that time where concerns about competition disrupted a negotiated deal for Do Not Track. Procedural matters are key for how standard-setting can support or inhibit competition and this experience particularly elucidates the different values transparency has and the role policymakers can play in the course of techno-policy standardization.

[Individuals vs organizations in standard-setting process](#) recounts the tradition of individual participation in technical standard-setting and illuminates competing views of standard-setting as stakeholder-balancing or technocratic and the individual's role as representative or expert. This is informed by the engineering ethos and

autonomous role of individual engineers laid out in Chapter 2 while responding to the research question of Chapter 1 about how collaborative governance can be inclusive and focus on solving problems.

[Who participates and why it matters](#) collects findings and analysis, both qualitative and quantitative, of who participates in technical standard-setting processes in general and the Do Not Track standardization in particular. Participation speaks to access and legitimacy of techno-policy standard-setting processes (as laid out in Chapter 1) but also provides context to who the designers of Internet protocols are and how their particular views will affect the implemented designs (as raised in Chapter 2).

[How participants see privacy](#) collects what participants in technical standard-setting processes think about privacy itself. Chapter 3 noted the ongoing and productive contestation of the concept of privacy and this section details the conceptions of privacy shared by participants along several different dimensions. The research question raised in Chapter 2 asks how these individuals' views of privacy affect the privacy of Internet users. This section explores how participants think about diversity in others' views and specific cases of considering the privacy of differently situated others, from their own children to Internet users at large.

Based on these themes, I see opportunities and challenges for the larger project of supporting privacy through multistakeholder technical standard-setting processes. [Towards integration](#) makes those connections and describes an argument for more nuanced integration of process, work and expertise.

5.2 How standard-setting accommodates, succeeds and fails

In considering the future of multistakeholderism for tech policy, I asked:¹¹¹ What are the impacts of techno-policy standards-setting processes on resolving public policy disputes for the Internet? How can we establish relative success and failure and what conditions affect those outcomes? Here I share some of how participants in technical standard-setting and the standardization debate over Do Not Track describe their experience with the process and what made it successful or not.

People I spoke with distinguished emphatically between participants acting in good faith and bad faith, as well as participants who were more difficult or easier to work with, all of which are orthogonal from actually agreeing on substantive issues. Interpersonal animosity has a significant impact on participation and interpretation of process, with its effects seen and felt in different degrees. As a result, open processes on questions of public policy values or on any questions that handle disputed topics must accommodate diverse perspectives and a variety of tactics from participants.

Success and factors for success can be examined in each of several stages of a multistakeholder standard-setting process: in incentivizing; in convening, communicating and learning; in agreeing; in implementing; and, in using. At each stage, participants may have different criteria for success and success and failure may include impacts of the process in other policy settings, not just the room where a standard is being debated.

5.2.1 Good faith vs bad faith Participants in standard-setting can identify good faith disagreements, even on topics that were fairly controversial: over privacy, permissions, etc..

I should say also that I don't mean to paint anybody in a bad light. I think that everybody in that debate was acting in good faith and had good reasons for what they were thinking.

There are lots of people that sit on standards bodies and they all come from different points of view. I think they're all doing good work and I think they all have best intentions, but we represent the user and the user agent, and ensuring that we have the flexibility to do what

¹¹¹See [Internet Standard-Setting and Multistakeholderism](#).

we need, and I think we have a pretty good track record of ensuring that we do do the right thing.

In contrast, an ad industry representative described industry participation in the Do Not Track standards process and NTIA-led multistakeholder processes as more calculating, using language emphasizing bad faith:

it was about as Machiavellian as you would think. It would be as backroom, smoking-cigars-in-a-steakhouse as you would think [...] and this is true in Washington. This happened with the multistakeholder process – and it’s naive to think otherwise – that people agreed beforehand who was going to be good cop, who would be bad cop, who would raise what points so it wasn’t always one entity; companies would agree. And so you were sitting in the room assuming good faith and everyone’s there to share the same goal, and that was not occurring.

While this is among the more vivid descriptions, this particular participant actually identifies bad faith behavior narrowly. Communicating elsewhere about how to organize participation (who will say what) or not sharing all the same goals might not be considered bad faith behavior by some participants. Concerns about bad faith may go further, as described below.

Good faith is explicitly identified as orthogonal to agreement on goals or outcomes (because why else would you need to describe someone’s faith as good); e.g. “did not agree a lot with what they did, but they were thoughtful and fair and honest in the room.” There can be similar positive evaluations of not just fair spiritedness but also taking reasonable or supported positions (which, again, others may disagree with):

he had arguments why it’s expensive. And then you can argue whether you say yes or no to the argument, but it was a substantiated concern. It wasn’t just saying, “I don’t like it, and my business will go down the drain, and I will go bankrupt, and whatever: the whole ecosystem will collapse,” these kind of statements, but he usually had a sound argument why a certain proposal was not in the interest of his company. And he could have basically just disrupted the process, and naturally he fought for his views, which is perfectly fine, but in a substantiated matter. So that’s something I liked.

The value of honesty and the potential of an “honest broker” position is also frequently raised by participants. While honesty is generally appreciated in order to work out disagreements, an honest broker is identified a little differently, often as a neutral, external or go-between party (government actors are sometimes described this way) who can talk to both sides¹¹² or all sides in a dispute and give honest assessments of what compromises are possible.

Descriptions of bad faith can be more diverse.

Sometimes it’s a question of the quality of argument or reasoning, making unfounded statements without any expectation that they would be useful. Comments and arguments are described as “absurd,” “completely clueless” or “ridiculous.”

Related is the criticism of “giving speeches,” a metaphor about speaking to communicate commitment to a set of positions but not in an attempt to converse with other people in the room. In some cases these are mismatches in audience – a representative is instead signaling to people elsewhere that they are repeating the approved position. In a notable case referred to by a couple very different participants I spoke with, an advertising trade association representative read portions of a letter (and pasted sections of it into the minutes) about their categorical opposition to Do Not Track and the W3C process, interspersed into a technical discussion of unlinkability.¹¹³ Language like “grandstanding” or “talking points” is used similarly.

And as long as we’re talking about advertising, I think there were certain particular advocates who wanted to grandstand about the evils of advertising as well as some industry people who also grandstanded because they were late to the process perhaps. And given the amount of travel required for those meetings, that tended to bother me because I would feel like I have a bunch of needy family members thousands of miles away, I didn’t necessarily come to hear you give this speech for several hours, right? I think I may not have been the only person who felt that way about either side, right? I think there were times when we were working more effectively to try to get things done and some of the speech making and the either anti- or pro-advertising stuff was not helpful.

Distinct concerns are about those trying to be disruptive to conversation altogether.

¹¹²See [Stakeholder groups: counting sides](#) in the section on participation.

¹¹³Minutes from October 2012 and press release re: [open letter from DAA to W3C leadership](#).

And then what I found also interesting, there were people specifically sent to disrupt the process [...] Sometimes there are contentious issues with different opinions, and that's something you can manage, but it's hard to manage people who just disrupt the room, shouting this and that and you're creating turmoil. That's an interesting challenge.

I think both sides have been, frankly, pretty ridiculous, just some of the behavior. I mean, I think you were at Microsoft, one of those breakout sessions where literally people had to be pushed back, that wasn't the privacy advocates doing that. <laughs> I'm not naming names anywhere, right, but that to me was just like, are you kidding me? This is like a New York thug here trying to, like, bounce on some people? I was shocked at some of that.

Participants also describe behavior as subversive of the process without being as directly disruptive, as in trying to delay decisions or discussions procedurally.

the endless delays, you know there were plenty of cases where it was perfectly clear this was a delaying tactic and we were going to spend three months handling a formal appeal or a formal objection or an appeal of something, right, and the outcome was going to be predictably that, no, this was a decently balanced decision by the working group and the chairs and it should stand, meanwhile we're three months later, and they pulled that handle multiple times, it was getting frustrating, you wanted to be able to say fuck it, guys, stop playing delaying tactics, we're just going with this decision, no, we're not going to hear your formal appeal, or your formal objection, or your appeal, but each time we said, okay, fine, we hear your formal objection

This particular description of appeals in Do Not Track is interesting as I believe the Formal Objection process, a W3C procedural step that can be applied to any decision, was only actually completed once. But it could be participants recall objections and appeals more generally, which were relatively numerous.

Others described slowing things down as an explicit and intentional goal that they thought was just a benefit to a more considered or acceptable outcome.

But I think over time, you've got to remember this was like a five-year process, so I think your initial goal is do no harm, let's get engaged, let's figure out what's going on here, let's put the brakes on this so

we can understand it, and then we can come back with considered opinions on what some options may be that we could actually live up to.

5.2.2 Animosity While sometimes aggression is identified as intentional disruption done in bad faith, it's also described as a separate phenomenon that arises from heated conflict. This theme comes up with standard-setting in general, but it's especially prominent in the discussions of Do Not Track, which was notably heated and antagonistic.

Animosity is typically defined as ill-will that involves taking action based on that hostility. That animosity arises is perhaps not a novel research finding: standardization of Do Not Track involved people with dramatically different backgrounds, representing conflicting interests and competing financial models, and without long-term experience working together in a shared community. Longer-term, regular participation and community development is described as one aid to lessen conflict in technical standard-setting more generally. While not unexpected, it is useful to note some of the effects that the level of acrimony had on participation and on the process itself, and how those effects varied.

5.2.2.1 Difficult people Participants being “difficult” is often a property identified about the people themselves, as separate from working in bad faith, productivity, or supportiveness of the process or its goals. That people in technical standard-setting processes can be difficult is generally known, as in this brief description from Bray (2012), which I think rings true:

Standards-making is a boring, bureaucratic, unpleasant process, infested by difficult people and psychopathic institutions. Some “standards” turn out to be useful; most are ignored; some are actively harmful. And looking at which organization the standard came from turns out to not be very useful in predicting what's going to happen.

The idea of “difficult people” comes up regularly among people I spoke with, with language like “prickly,” “not terribly pleasant” or “difficult to work with personally.” While difficult-ness is not directly attributed by interviewees for negative outcomes, it is sometimes considered a distraction or it's noted that it “didn't always help.” While this study doesn't have sufficient depth on this particular point, it would be worth exploring how this commonly accepted quality among some technical standard-setting participants may be discouraging or disruptive.

Many of the behavioral characteristics described here appear to be gendered; the people specifically identified as difficult were most often men. Throughout many open source software projects there has been a push towards codes of conduct and W3C has had groups working on procedures for Positive Work Environment since at least 2007;¹¹⁴ those efforts have also faced pushback which has typically demonstrated the presence of discouraging and antagonistic behavior and the need for more welcoming environments.

5.2.2.2 Toxicity and personal attacks Beyond simple difficulty, some participants explicitly identify toxicity or personal hostility as discouraging participation and leading formerly engaged participants to exit the group altogether.

Any one of the times [A] and [B] got into a screaming match on the mailing list. How do you deal with that? Great, your standards body has turned into a flame war. I never had a good answer for how to handle that, but I knew that it was highly destructive. It silenced some of the members. We lost [C]. I mean, there were just wonderful people who no longer wanted to be near this toxic environment and I couldn't blame them.

[...]

I don't know how to solve Gamergate. I don't know how to solve people deliberately being mean to other people to try to get their way. That may not be what you thought of, when we were asking about fairness, but it was fundamentally unfair. It was silencing people by being obnoxious, and it was effective.

This participant in the DNT process identifies a “toxic environment” as a particular issue of fairness, or of procedural legitimacy, that we might not traditionally identify. Similarly, not maintaining civility is identified as a failure from leadership to protect a participant who described feeling pushed out.

you know, it did contribute to me leaving. Which, as I say, again, I think like one of the responsibilities of a chair in a working group like this, especially when it's going to deal with tricky policy-esque issues where there might not be consensus, it's to a minimum keep the place

¹¹⁴Positive Work Environment statement of principles [appears to date to June 2007](#) and a more formal version was published this year [Siegman, Li, and Cannon \(2020\)](#).

civil. Right? That doesn't seem like – these days I guess maybe that is too much to ask, but at least at the time it didn't seem like too much to ask.¹¹⁵ And it really bothered me that the working group chairs didn't seem to view that as a priority.

I feel that part of this critique is directed towards me as helping to organize and manage the Working Group, and I take it to heart as a valid and important criticism. While we had some offline conversations with individuals about civility and chairs of the group had occasional guidance on civil and constructive behavior on calls and mailing lists, retrospectively I can see how little preparation there was and how few controls were in force.

W3C and IETF have had policies in place to occasionally warn individuals and restrict participation in egregious cases, but they were designed to be used very infrequently, assuming a self-governing and mostly homogeneous set of professionals in a tight-knit field. These policies seem woefully out-of-date today. Both W3C and IETF have initiated some processes more recently to better handle violations of professional conduct, but it's still often a struggle and controversy when they're employed.

It may not be settled what conditions of civility are expected or what norms from other settings should be used. While many identify antagonism, conflict, personal attacks and incivility as common and disruptive, people view the degree and importance in widely different ways. Some called it “no different than in any other workplace in a way” or that it was remarkably civil despite having conflicts and disagreements. Some identified strict process – about speaker queues, limiting speaking and threatening to remove troublesome participants – as reasonably successful at managing disruption.

It was a fair process. I thought given the task that we had, I thought Aleecia did an outstanding job of just trying to keep it sane. We have a ridiculous amount of conflict in that group and it's not like other standards working groups where two competing implementations might have different ideas of how something might be done. This is a group where a significant portion of the participants were suing each other in court on different cases and they're on complete opposite ends of the spectrum. No desire to compromise at all. And yet we still had pretty civil meetings. So from that perspective it was fine.

¹¹⁵Interview was in 2018, but this is referring primarily to 2011-2013.

Reconciling this range of perspectives about conflict and civility is challenging for me. At first I thought it might just be an individual's own behavior – if you're more direct or abrasive yourself, then you might not be affected by toxicity around you – or about how personal the experience was – if you felt directly targeted, then you'd care, but if it was directed at others it might not matter – but neither of those heuristics fully explain the variations I see.¹¹⁶

Instead, it seems that some individuals (myself certainly included) tend to be deeply affected by attacks, aggression or animosity in a way that chills, disturbs or discourages participation; at another end of that spectrum, some other individuals find that roughness to be a common or integral part of work or politics, maybe it's even enjoyable or seen as active and direct. The source of that difference is psychological, and well beyond the scope here. But process that accommodates participation of the latter kind will tend to discourage or disrupt participation of the former kind. Looking back, I could have stepped forward to encourage aggressive – if manual and case-by-case – enforcement of basic rules for participation. Looking forward, what would a process with a modern code of conduct, an active commitment to maintaining constructive conversation and easy-to-use tools for moderating and blocking participants be like?

5.2.3 “Here's the process. Follow the process.” Several people I talked with emphasized the importance of a mechanical, regular application of decision-making process to make progress, to settle questions and address objections, without which there could be no end to any debate among entrenched parties. Partly this comes up as a reaction to identified tactics of delaying: that some would prefer no progress to be made because of the potential effects on their business or the external effects of the debate remaining ongoing. Partly it's a reaction to how slow standard-setting processes can be generally and a frustration with the time, cost and delay involved. And finally it's described as a characteristic of fairness, a way to ensure that all concerns are addressed without having to rely on either the good faith or the impartiality of anyone in a contentious debate.

with advertisers saying that advertising is as American as apple pie and they want the bug report submitted. Yeah, that's back to your fairness, right? “Oh, okay. You'd like to blow up the whole process? You'd like to exempt all advertising? Great. Here's how you file an

¹¹⁶Several heuristics might still be partial explanations, though, including also gender, professional background and cultural attitudes.

issue. This is the process to do that. We'll take it up in turn." Just straight up. "Here's the process. Follow the process."

I learned a lot about it as I went along, but over time I kind of – I developed a lot of respect for Roy and his kind of approach to it, which is very brass tacks, and here's the process, and process in place for a reason, and obviously everyone else is trying to hack on what they had previously brought to the discussion to the process to subvert it.

we're going to crank the handle, we're going to record decisions, we're going to set deadlines, we've got somebody operating the process now, that was the other thing I think really helped move things along.

Systematization may be an essential characteristic of process itself – if it's not systematic and consistently applied, then there may instead be a lack of process. The systematic administrative quality of a process is credited with fairness (to procedural legitimacy, again) but also with other positive outcomes in terms of reaching resolutions. And in all of these cases, participants have identified the systematic nature of the process as essential for working with a heterogeneous group including those antagonistic to any outcome.

Others felt uncertain about the details of the process, and that uncertainty is described by advertising industry participants as a reason for entrenchment or objection.

But conceptions of authority, escalation, you know, ultimate decision factors, those were quite – at least not with a degree confidence, were not well understood. [...] we would just say, "Okay, do you understand this?" And other people would like, "Well, I think I do, but not really." And if we ultimately disagree, what happens? That's where people just did not [feel certain] [...] It didn't feel like we were falling back on, "Hey, we've done this process for twenty years. This is exactly how it works." It didn't feel like we were in that situation. It felt more like a, "Hey, we've never been in this contentious of a situation before, so we're kind of building this process on the fly."

Rather than either supporting a set of procedures or objecting to them as unfair, one can also be simply uncertain about a process or its effects. That this uncertainty is tied by this participant directly to questions of escalation is likely relevant: W3C has documented process for escalation, appeals and objections to

group decisions, but there is a sense that those procedures should be only rarely invoked in a consensus process, which might contribute to a sense of uncertainty about them. That processes operate with both norms as well as formal rules is not unusual, but it may be that contention generally causes a push past norms to more formal rules when escalating objections. (The function and dysfunction of the US Senate in our state of political polarization seems one topical example.) Uncertainty about the formal process and the informal norms may also be tied to a lack of previous experience with technical standard-setting, which was common in Do Not Track.¹¹⁷

Well, there were hundreds of proposals and amendments and, you know, questions to be— I know there wasn't really voting per se. I know we joked about it sometimes, but, you know, we would hum about proposals, whether we agreed with them or not, because it was supposed to be consensus. And a quick aside, I think the process was never actually clear to a lot of us. You know, to me consensus means everybody agrees. We all say whether we agree or not. I think there were a lot of questions about what consensus meant on certain proposals. [...] this is what I've always told my teams. We can all disagree about substance. But when the process is broken, you can't defend that. And I think there were times when the process was broken. And when the process is broken, the substance doesn't matter. If people don't feel that the process is fair, then we all can't agree that water is wet. Somebody will disagree to that.

This seems like a key point on the connection between fairness and agreement, or between procedural and substantive legitimacy. Where there is a lack of trust (again, whether that's uncertainty or a particular concern about unfairness) in a process, then disagreement on questions of substance is even more likely, even on the simplest of questions.

While some describe the TPWG process as simple and systematic, others found it unclear or uncertain. For both assessments, though, there seems an aligned interest in clarity and systematization. Precision in a process might improve trust in a contentious process, and it might also improve fairness and progress even when groups remain highly contentious. Is process, then, just an unalloyed good?

¹¹⁷See "Standard-setting organizations as social networks" in the section on participation for more on this topic.

It can be commonplace to complain about bureaucracy, tedium or formality in standard-setting processes, that it would be faster and easier if everything were just informal and quick, like decisions made inside a company or on a software project. Surely there is some balance to be had there. But notably those complaints about bureaucracy were not emphasized by the people I spoke with. While concerns about the slowness and amount of time spent are raised – as questions both of fairness and of efficiency – those are more often concerns about not making decisions, rather than having a process to systematically address concerns and resolve decisions.

5.2.4 Different areas for potential success or failure Most participants I spoke with would say that the attempt to standardize Do Not Track was a failure, or at least wasn't the success they had hoped for. But at the same time, it's very common for participants to identify different kinds of outcomes as successes or potential successes, or moments where they believe a change was necessary for success by their criteria.

Consider a kind of progressive timeline model for creating a technical standard. First, there must be incentives in place for organizations to have reason to consider a change or a problem that motivates a standard. Next, you have to actually get the stakeholders, particularly the stakeholders who might use the standard one day, into a room to work on the common project. They need to talk with one another, and hopefully learn and understand their different positions better than before. Once they're talking, success requires some level of agreement, involving compromise or some form of consensus. Agreement is only a pre-condition, though, to actually building something: implementing the standard in software and services and putting it out in the world. And finally, use is also not deterministic, there have to be people who use the implemented technology and some consideration of whether that use addressed the initial motivated concern.

We might visualize it as similar to a software development lifecycle model, as in this diagram. There are some similar concepts: defining the requirements or needs of a problem that needs to be addressed, settling on a solution, implementing and using it or testing it out. A very traditional waterfall model would fit the initial linear idea, though it's probably even less realistic in the case of standards, where implementation and use both drive and are driven by standardization, but it's a starting point.

This model doesn't identify any step as especially important or especially challenging, but it helps us to understand the variations in what participants

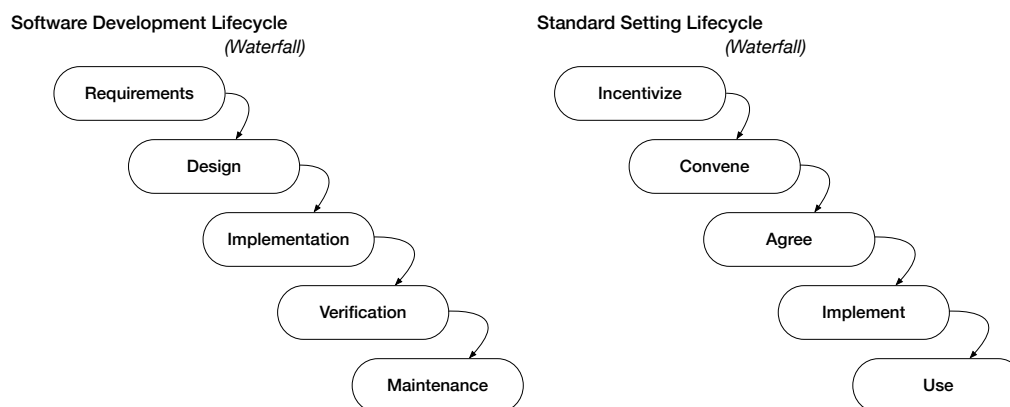


Figure 14: The traditional waterfall model of the software development lifecycle, with a somewhat analogous set of steps for developing a technical standard.

identify as both successes and failures or the reasons to which they attribute success or failure.

5.2.4.1 Incentivize For any technical standard, there needs to be a direct incentive for participation, both for the development and ultimately in implementing and adopting the new protocol. Incentives are necessary because there are significant costs to each stage of the process: it costs time and money to follow, attend meetings, negotiate alternatives, as well as to develop new or updated software or make procedural changes to meet a new standard.

For standards addressing integrated technology-policy concerns, incentives are just as necessary, but we might identify a broader range of incentives beyond the more typical direct market needs, like deferring or avoiding regulation, facilitating compliance with regulation, and addressing social or political concerns.

Incentivization also has an influence throughout the lifecycle (in contrast to the abstract waterfall diagram described in the overview): if there's a threat of regulation early on, that might be enough to bring companies to the table to explore an alternative, but if external changes make legislation unlikely, that can reduce pressure to continue participating (and continue paying those costs) or to implement a completed standard.

look, once the Republicans took the house, I mean, that was a lot of the momentum gone

Incentives are often explicitly tied to the step of convening, of “people in the room,” but also just as an analogy to moving forward with any step of a voluntary, multistakeholder process.

It’s nice to put people in the room, but, I mean, unless there’s a reason for them to be there, unless there’s a reason for a company to say, “Yes, I am willing to make this change that will cost my company money,” then why would they do that?

That some form of a pressure is a necessity for multistakeholder success and for privacy in particular was raised by some participants directly in connection to the Obama administration’s proposals for multistakeholder processes convened by NTIA – and, emphatically, that the multistakeholder process was supposed to accompany privacy legislation which itself would be an incentive.

There were initiatives that the president pursued concerning global interoperability of privacy frameworks. There were initiatives involving proposals for national legislation. There were initiatives involving pulling more technologists and other smart folks into the federal government to try to get greater privacy and technical expertise at agencies like the Federal Trade Commission and the White House and others. [...] the multistakeholder engagements that the president contemplated as being, one, potentially standalone initiatives and, two, as being complementary to the idea of a national privacy legislative strategy that would provide incentives for folks to participate in those sort of engagements and for those engagements to be more worthwhile.

Some participants in the Do Not Track process who were less familiar or had less background in policy identified regulatory pressure as important, but only “in hindsight,” recognizing its importance after the process had stopped or waned.

While identifying legislation or regulatory movement is commonly identified as an incentive for industry participation in self-regulation or in multistakeholder negotiations more generally, it is not the only relevant kind of pressure. Many participants identify blocking – of cookies, of tracking, of advertising, of various kinds – as relevant incentives that moved or could move negotiations forward.

There’s a huge difference in life between a threat and a credible threat so there would be a theoretical threat that the browsers could do this

but there was momentum behind the idea that spring. Going into the Apple meeting, there was enough solidarity among the browsers that I thought we had a really good chance of doing it and the third-party people were treating it much more seriously than they had previously.

As described in the [handoff analysis of Do Not Track](#), technical measures (or the credible threat thereof) can be actions taken to create a handoff, a shift between different paradigms. The threat of more extensive blocking measures of online tracking mechanisms by browser vendors was used as an incentive to shift from traditional notice and choice towards a cooperative Do Not Track agreement.

Some identify more interest in Do Not Track now¹¹⁸ than during its time of development and discussion because of recent increases in both legal pressure and technical mechanisms. The European General Data Protection Regulation (GDPR), which puts stricter and more consistent requirements on companies handling personal data of European residents, puts added pressure on gathering affirmative consent for many kinds of data collection and DNT has been suggested as a way to more efficiently communicate that consent (O'Neill 2018). The California Consumer Privacy Act (CCPA) provides California residents with the right to opt out of sale of information collected about them, which may include browser-based tools like a Do Not Track setting. Browser blocking mechanisms, including blocking or limiting access to cookies or blocking requests altogether, have continued to develop over time, as part of the new arms race paradigm. Legal requirements and technical measures both may exert the kind of pressure described as an incentive for participation.

5.2.4.2 Convene One view of standard-setting is that it is essentially about gathering people together. This is sometimes epitomized by the language of those most long-term involved with standard-setting, and it was much of my experience being employed as staff in a standard-setting body. Under this view, there is a value and a chance of success just in getting the different relevant parties involved talking about a well-scoped issue. It also implies a certain sense of neutrality in that convening is a priority in the sense of getting participants to identify the particular outcomes themselves through the convened process. The restaurant-with-tables perspective is one relevant metaphor.¹¹⁹ Convening has also been a key tool of

¹¹⁸Circa 2020.

¹¹⁹As described in the chapter on [Internet Standard-Setting and Multistakeholder Governance](#), citing Bruant (2013) and a description of W3C.

FTC and other agencies in pursuing new governance approaches rather than more traditional, formal rule-making.¹²⁰

Convening is inherently tied to the questions (discussed above) of incentivizing participation.

And I can't get anything done if nobody feels any pressure. I mean, I can get something done, but people aren't necessarily going to commit to it or aren't going to put a lot of time into it. We still may do it, because we might think it's useful as guidance and it'll be useful when people are ready for it. But it may not have as much impact. It may be an interesting thought piece. And we do that. I think lots of groups primarily do that. [...] what W3C does when things are hitting on all cylinders is engage people who see a need any time for various reasons. They want legislation, they don't want legislation. They need a standard, they need something to work, to build, to solve a problem, whatever the complicated motivations are, people are ready to show up and then it's a matter of choosing the right stakeholders who credibly represent, but are ready to work and getting them in a room. And when you do that, smart people come up with good stuff. They're incentivized [sic] to get to an end point.

There is a wonderful optimism (that honestly, I share) in a motto of "smart people come up with good stuff" that exemplifies this convening mindset. But I include the longer quote here to show the necessary pre-condition of incentives.

At the same time, convening itself is sometimes attributed with making subsequent agreement more challenging: getting too many diverse perspectives may encourage more opposition or otherwise make it difficult to settle on a particular solution: "it's much harder for making a deal." What some identify as success in convening disparate parties is identified by others as a roadblock to successful agreements.

Getting the parties to talk is often cited as having its own benefits, even separate from whether they agree upon a standard. Learning about different stakeholder positions can be beneficial to better understand underlying conflicts, understanding detailed questions can lead to more fruitful conversations, and developing working relationships with other parties can promote improved handling of future conflicts.

¹²⁰See, for example, [Cohen \(2012\)](#) and [Bamberger and Mulligan \(2010\)](#).

Learning is identified as helpful both for advocates to understand business and technical practices and for people within industry to recognize consumer or civil society concerns in more detail. (There are many, many more of these quotes than I'll include here. Maybe it's mentioned so often and at such length because it's a feel-good conclusion in an area that otherwise sounds so contentious. Or maybe it's because it's a genuinely distinct and important benefit that many people identified during and after the fact.)

I think I've gotten a better understanding of different people's perspectives and what it is that they think is important. I think in the job that I have, which largely is bringing input from outside the company into the company so that it's a part of our decision-making process, I think being a part of the working group and hearing the way that people talk about different issues has been really helpful. I think it gives me a good sense, and by extension gives our company a good sense, of the way that people address issues and the things that people are going to be concerned about and those sorts of things. [...] what are the big picture priorities for somebody that's a privacy advocate in our group? What are the things that they are worried about us doing? What are the things that they are not worried about us doing? One of the things that I think is the case is I think we see a narrowing understanding gap on both sides. So, I think one function of the working group, which is not necessarily the core thing that we're all gathered to do, but I think it's at a consequence, is letting me understand what their priorities are and also hopefully giving other folks in the group a better understanding of the way that we approach privacy and the way that we approach information.

I was very concerned that the policymakers fundamentally didn't understand the issues. Now, it turns out the technologists in a lot of cases didn't either. I think one of the real pluses to the Do Not Track process was that a number of people learned a lot of stuff from one another. So people in industry might have a glimpse of their particular part of the picture but not understand the rest of the ecosystem, and so even for the people writing code that did stuff, they got surprised. And I thought that [...] if the goal of DNT had been a fact-finding mission, that is actually useful. That part was good. But I was concerned that people were going to write laws without understanding

the underlying technology and would write things that were either technically impossible or just stupid.

That participants from all sectors identified learning about other stakeholders or technical details as helpful and frequently cited them as successes doesn't mean that communication or information was always easily distributed. Several people I spoke with identified "information asymmetry" as an ongoing challenge for those outside the ad tech industry; at times during the DNT process consumer advocates would request internal details on industry operations that would not be fulfilled.

5.2.4.3 Agree Getting to agreement, consensus or compromise depends on success in the previous questions on incentivizing and convening. Multiple people I spoke with referred to the concept of the best alternative to a negotiated agreement (BATNA (Fisher, Ury, and Patton 2011)) and how organizations (whether ad industry or consumer advocacy) would be influenced by what they believed they could 'get' outside the process. And free and productive conversation is identified as necessary to help find the set of compromises that would be acceptable, the zone of possible agreement, or ZOPA,¹²¹ (although some interviewees use alternative terms of art for this concept). But there are other challenges in reaching agreement when parties or individuals are entrenched or in finding a rough consensus from a larger group.

5.2.4.3.1 Barriers and getting to some consensus Barriers to consensus or compromise that are proposed by participants are diverse and interesting, if also speculative. Peer pressure within a group (say, of advertising industry executives, or civil society advocates) is cited as discouraging reaching out or making concessions towards a compromise. Discussions in public, or decisions that could be quickly reported by media, might "make people very cautious." Some individuals might have less personal incentive to make a compromise compared to what would benefit their employer or members,¹²² and it may be particularly challenging to run a consensus process where many are either opposed to any agreement or more focused on delay or uncertain about how the process does or should function. Entrenchment, or even the perceived entrenchment of others, could make individuals less likely to approach the process as a collaboration.

¹²¹See Sebenius (1983/ed) for one relevant description.

¹²²See *Individuals vs organizations*.

How consensus is defined and identified is distinctive and important for decision-making in standard-setting processes. Under some political definitions, consensus is unanimous agreement (or lack of any objection), and standard-setting bodies have typically pursued a “rough consensus” approach given the challenges of letting a single person veto any decision. For the Tracking Protection Working Group, that included a model of a “Call for Objections” when further discussion seemed unfruitful, where each option could be considered and participants were polled on what their objections were to each option and the chairs of the group would identify the option with the least strong substantiated objection to it. While the details of that process are not very frequently mentioned by participants I spoke with, it is sometimes referred to as a “forced consensus” or like “adjudication.”

So what we did in the [Tracking Protection] working group, we forced consensus. We had a process saying, okay, we put all the options on the table. People can raise concerns. And then [chairs of the WG] basically took the option which had the least substantiated concerns, which sort of is a way to force consensus. And then what happened is basically— so on the technical side, everything converged. We have a standard. It’s sort of, in theory, successful.

I think we tried to look at the HTML experience as we were modeling the new way, we were doing consensus through written texts.

That contentious decisions are made is not uncommon in the standard-setting space, as in this reference to the HTML Working Group. But there is some idea that for a voluntary standard, contentious decisions have to be made and recorded even if the outcome won’t be acceptable to everyone, but enough people still have to be willing that a standard could be meaningfully voluntarily adopted.

Eventually you land on consensus of a self-selecting subset of your stakeholders that for some reason sticks to their own process. And you then find out whether that subset overlaps with your implementers.

As this long-time standards participant notes, for a model of standards as successful if implemented, the final “subset” must contain enough implementers. For legitimacy evaluated through other means, it’s less clear or at least less defined what enough of a subset would be: if unanimity is unavailable, how much agreement must there be among how many of the interested parties to declare a consensus process successful?

5.2.4.3.2 Deal-making Many contrast to other types of negotiation and in particular, there are both positive and negative comparisons to “Washington deal-making.” Trust (or lack thereof) among participants and the relative ease of smaller groups in closed door settings are raised, and connected to the anti-trust and transparency discussions in [Competition and standard-setting](#).

Contrasts are drawn to negotiations from the federal legislative process, where, for example, groups with power, interest, expertise or diverging opinions have:

the ability to come to the table with one or a couple of speakers– one or a couple of representatives, and then hammering it out. And being able, then, hopefully, to deliver for their group. So that’s how things tend to work in Congress, when it’s a hard negotiation.

But some are explicit about the incompatibility of that approach with a consensus process, because of the lack of transparency and the lack of trust.

It doesn’t work like that, and I think that was really a period where we wasted a lot of valuable time in terms of solving the problem. It also put the whole Do Not Track process at risk, because the real Washington deal-making was not done in the room. It was done in a parallel process in a table with four or five individuals. Nobody knew what’s really happening there. I got some [...] output, but I didn’t get input, and it was also based on a very loose promise. There was not a lot of trust.

Transparency and participation in how deals are made may have effects (positive or negative) on the likelihood of reaching an agreement and can separately have an effect on the success of an agreement having sufficient legitimacy or stability.

5.2.4.4 Implement In contrast (again) to the idealized waterfall model, setting technical standards is typically driven in part by implementations: implementation experience is a necessity and there’s often no incentive to standardize until there are some rough implementations to talk about. In the case of Do Not Track, implementation by browsers came early as a way of kick-starting the idea of a user preference about online tracking, but implementations by large online trackers was lagging and uncommon.

Early implementation of sending a Do Not Track via a user preference is identified as an encouraging sign:

I feel like I was hopeful at the beginning of the process. I mean, the companies had all agreed to put the browser instruction in the browser. So, that was a good sign, and it seemed likely that something had to happen because of that. Not just that it just might sit there forever being useless. So, I think I was hopeful.

Voluntary implementation by companies in online advertising was a particular barrier given the potential revenue impacts, and several people I spoke with identified that as a basic control that ad industry had.

all the stuff that people really care about happens in the back engine room, and so ultimately doing something without the servers sort of being part of it from the beginning is gonna be difficult. They have the ultimate leverage, right, until law tells them to do otherwise or until there's a harm so gross that they as human beings have to do it, right? That's the challenge on this one.

Implementability can also be identified as an indicator of substantive legitimacy or success in identifying the right solution:

making sure that companies are actually able to implement this. I think there's one vision of the right result that – and I'm sure people have said this, "if companies go out of business because of Do Not Track, that's okay." I don't subscribe to that view. I think there should be a way to do this without putting people out of business, without fundamentally changing the ecosystem. And so, maybe it depends – I think part of what the right result is something that people are going to voluntarily do.

Often going hand in hand with the perspective of implementations as impact (and therefore as ultimate success criterion), is a sense that there has to be some analogous kind of adoption to affirm the legitimacy of the outcome.

Yeah, people were kind of focused on who can practically implement, who can use their voice to make sure that this is credible, and then there were other considerations there too, but I think those two things are pretty important things.

Some describe it as the overlap between something that is built but also helpful or valuable:

get people to agree on a core set of things that maybe could work and then hopefully build something that is in the intersection of what people are doing and are willing to implement on the one hand and what actually makes a useful difference for users on the other and that the more privacy-leaning parts of the parties in the conversation could actually agree to.

Implementations can also have impacts without standardization, and can provide experience for other implementations in other areas.

Hey, look, it ended up leading to the mobile OS systems developing their Do-Not-Track-like tools, which, again, probably wouldn't have happened. Almost certainly wouldn't have happened. Other platforms have the same things. When I looked at Smart TVs, a lot of them had mobile OS-type of controls to limit third or I guess fourth-party, however you want to look at it, data collection and using rotating app identifiers and so, I mean, it kind of just helped put pressure on industry in various ways that I think was productive if nothing else.

This participant notes that mobile operating systems have tracking limitation user preferences, and the primary operating systems are developed by companies which also develop prominent browsers. Changes to the operating system (what we might identify as another platform, along with the Web) don't require the same level of cooperation from app developers, but these settings seem to be directly influenced by the model of Do Not Track: a simple binary opt-out request that a user makes in a central device location. Success through implementation could happen inside or outside the direct standardization process.

5.2.4.5 Use Even technology that is developed – coded, tested, deployed, etc. – makes little impact without users. Many describe the ultimate lack of success of Do Not Track is that users don't have any reliable functionality: you can't flip a switch on your browser or device and opt out of online behavioral tracking for advertising or other purposes. Despite the lack of server-side adoption and meaningful functionality, sending the Do Not Track header from users' browsers was actually quite common (at times reported at over 10% or even over 20% of

visitors)¹²³. Failure for wide-scale adoption may even be attributed to use being too high.

While different participants described their success criteria differently, many included a theme of having a DNT signal that a user could select and have a meaningful outcome. Some were explicit in hoping it would be adopted by a small portion of people, in the hope that that would make a significant opt-out more acceptable to sites and browsers that implemented it.

so my version of successful would have been 2 percent had enabled DNT and there had been a standard published from W3C with adoption by a handful of the top websites. So what we saw was user adoption was way higher than my 2 percent hope. It's like, 17, right? But that adoption by companies was extraordinarily thin.

That too many users might enable Do Not Track, or just the uncertainty, might change the financial incentive passed through a company hierarchy:

And it's the uncertainty that killed the adoption because people just don't know what they are going to do, you know. They don't know whether this is a feature that will result in, you know, a million dollars benefit to their customers in exchange for maybe a \$2 million dollar loss on the revenue side, yeah, that's okay. You know, or is it going to be a \$50,000 dollar benefit for the customers on a \$5 billion dollar loss on the revenue side, like, eh, that's not going to happen, you know? Because that ultimately is the discussion you have with the CEO when you get to that frame when you're going to deploy it internally.

One driver of high DNT usage statistics was a decision from Microsoft to turn on DNT by default (or within the bundle of settings that users could confirm at once) for their Internet Explorer browser. While participants I spoke with had different perceptions and explanations about Microsoft's decision, one common thread about its impact is that it could or would make for usage numbers that would be unacceptably high and therefore discourage adoption by industry.

Another kind of use, or re-use, is raised by some participants I spoke with: the re-use of Do Not Track, the concept or the technology or the specifications

¹²³Numbers that were collected and reported varied a lot, by browser, by site, etc.; one survey run by privacy-focused DuckDuckGo reported 23% of US adults in 2018 said they had turned it on ([“The ‘Do Not Track’ Setting Doesn’t Stop You from Being Tracked” 2019](#)).

or the discussion, in other settings for enabling user preferences. Most directly might be the California Consumer Privacy Act (CCPA), which I'm told was very directly influenced by DNT specifications and mailing list discussions. And interviewees later in my process mention the related ballot proposition, Proposition 24, apparently approved in the 2020 election as even more directly related, that it “doubles down [...] talks about plug-ins, web browser settings, and operating system settings.” It seems likely that the newer proposition would more directly support legal requirements for respecting standardized preferences for communicating opting out of data sharing (Edelman 2020) and a proposed Global Privacy Control closely follows previous DNT specifications.¹²⁴

This possibility of re-use of standards might again return us to the low fidelity of the waterfall model of software development. Indeed, the use stage of development may lead to testing, learning, and iteration on new cycles of incentivizing, convening, agreeing and implementing, in the same venues or in entirely new ones.

5.2.5 Conclusions for success throughout a process Success and failure can be evaluated within providing incentives, convening the right stakeholders, getting to agreement, implementing a standard and using it in the wild. But in each area, participants also identify ways that a process can affect the results at other stages: convening more broadly might make it harder to get agreement or convening a smaller closed-door group might affect the legitimacy of an agreement. Implementation and use might be pragmatic necessities for making an impact, but their impacts can at times also discourage agreement or can seed the ideas for future multistakeholder convenings. These factors affecting success can be seen in more detail in a particular series of events related to competition and transparency during the Do Not Track process, described in the next section.

¹²⁴Unofficial draft: <https://globalprivacycontrol.github.io/gpc-spec/>

5.3 Competition and standard-setting

Competition is a potential concern in any standard-setting project, because it could be used for collusion of some players against others. This concern comes up with DNT and privacy in a few ways:

- 1) a concern that smaller players in the ad industry will be relatively harmed compared to the larger players (or in web publishing, or between companies that had first-party interactions vs those who were solely third-party, this argument can be a little fluid);
- 2) a concern that browser vendors could have anti-competitive liability for agreeing on some set of limitations on their privacy tools related to a DNT compromise, and;
- 3) a concern that publishers or ad providers might inhibit competition by approving or disapproving lists of browsers.

But some informants with standard-setting experience also explicitly note standard-setting as valuable for competition: it lets them compete on other things because there will be interoperability, rather than a browser-wars situation of incompatibility.¹²⁵ Procedural matters are especially emphasized here: due process, transparency, dispute resolution. And in the case of an anti-trust concern arising over a DNT compromise in the spring of 2013, we can get particular insight into the different roles that transparency may play in the effectiveness and legitimacy of governance processes and how policymakers contribute.

Competition concerns and due process in standard-setting is historically especially related to intellectual property and patent encumbrance. That comes up remarkably infrequently in my interviews with standard-setting participants and was rarely mentioned in the context of Do Not Track. That was a surprise to me, since I prompted interviewees regarding legal considerations and because there was a documented issue of a patent that might inhibit use of expressed privacy preferences in Do Not Track, where a separate patent group was formed to investigate and resolve the issue. We might take that as evidence that the patent didn't ultimately play a significant role in DNT negotiations or implementations,

¹²⁵See “[The Web, Recommendations and Living Standards](#)” in Chapter 1 for a brief description of the “browser wars” and incompatibility of features between browsers and websites.

or simply that this study does not provide any further insight into the question of patent encumbrance and the effect on standard-setting.

5.3.1 The room where it happens Open door versus closed door negotiating is a common debate about multistakeholder efforts in general and shows up regularly, with different positions, among my interviewees. Some who are more familiar with self-regulation processes note explicit advantages of relatively closed door processes, because it can encourage candor among participants or provide less pressure of how particular negotiating positions may be reported in the press (or otherwise) and because it can facilitate packages of negotiated compromises. W3C's standard-setting process, on the other hand, follows an increasingly open-door process, with public minuting of every meeting and publicly archived emails of conversations. That openness is also cited with advantages from some interviewees: there may be legitimacy advantages of open discussions in contrast to “smoke-filled back rooms” (some version of that cliché is used by people on both sides of this particular mini-debate), it provides increased access to those who may not otherwise be guaranteed a role in a smaller closed-door meeting.¹²⁶

Attempting to combine the benefits of both such approaches, Peter Swire¹²⁷ tried to both engage people in very regular and openly documented meetings via teleconference and face-to-face meetings while also facilitating closed-door negotiations among a smaller group, including senior players in the advertising industry. This is not entirely novel for W3C or other open standards processes in the sense that side conversations (another notable theme among interviewees) are not considered illegitimate or unexpected in a standard-setting process: *of course* people are always having multiple conversations with different individuals, groups and subgroups; the goal of the openly documented process is that the ultimate decisions will get reviewed, debated and made in that open venue, after much discussion has already happened in various private and public fora. That practice is sometimes extended even to meetings that everyone can attend; for example, IETF groups have a policy of confirming on the list even decisions that seemed to

¹²⁶Openness can refer both to access (who's able to be in the room) and procedural transparency (who can see the details of what happens in a room); these are often, but not always, aligned. Here, “open door” will refer more to the transparency dimension.

¹²⁷Swire, a well-known privacy scholar and former White House official, was recruited as a co-chair of the Tracking Protection Working Group in 2012, taking over from Aleecia McDonald.

have consensus at a synchronous or face-to-face meeting, and that concept comes up in other multistakeholder settings as well.

Swire and others believed that there was a workable compromise that emerged from those private, high-level discussions, that could subsequently be discussed and accepted by the larger Working Group involving advocates, policymakers and various sectors of industry. That negotiated proposal was documented in a brief form prior to the May 2013 face-to-face meeting: it would involve the DAA trade association making respecting user DNT signals as an advertising industry self-regulation requirement, along with some efforts by browsers to make it not too easy to turn on a Do Not Track signal.¹²⁸

Some interviewees found this negotiated agreement feasible for many parties – the closest that the group had reached to that situation. But some thought it would not receive acceptance from a larger swath of industry when more thoroughly reviewed.

The default is cookies will be set, you can be tracked. If you hit Do Not Track, then it's what we said. That agreement would be collusion. That they could not agree on. I think any antitrust lawyer would have told you they can't sit in a room and agree to that. [...] So if we could have gotten to an agreement that it was off by default and that the X percent of consumers who wanted to not be tracked and exercise choice, we could have gotten Do Not Track.

So I think we were real close. We were probably just a couple of weeks away from having a vote within the working group on a proposal that I know for sure I had said I would support [...] And I believe the industry would have supported it, but then it got pulled at the last minute.

There was a short period there where I thought we had a deal. The browsers were going to be tough enough on the advertisers and the privacy people were going to get enough of what they needed that I thought we had a deal and then it fell apart, as deals sometimes do.

¹²⁸The six point “Draft Framework” was documented in a short document in April: https://lists.w3.org/Archives/Public/public-tracking/2013Apr/att-0298/one_pager_framework_as_distributed.pdf

my conclusion was [...] some of the [advertising self-regulatory] groups didn't actually understand enough of what they were talking about. So they could think they had agreed to something and then their member companies would find out about it and be like, "no, we can't." So that was my conclusion of how that was going to die, was that as soon as it got to a wider audience.

It surprised me how often I heard that a widespread agreement on Do Not Track was simply impossible (because of business model impacts, or lack of trust among parties, or with the inadequacy of the forum for discussion) but also that an agreement was basically settled on and would have been acceptable if not for a single hurdle in the Spring of 2013. Sometimes a person has expressed both of those views to me in a single conversation.

The hurdle in this case was a concern expressed by (at least) representatives of the Federal Trade Commission regarding the anti-trust implications of the negotiated compromise. Concerns had been expressed prior to the May 2013 face-to-face meeting; in the publicly archived minutes, Swire explicitly notes that he thinks anti-trust is not a concern with the proposal because of the general improvement in consumer welfare and choice through the adoption of DNT.¹²⁹ But a repetition of that issue in side conversations at the Sunnyvale meeting led to a private huddle of some of those participants – while the rest of the Working Group had an extended coffee and snack break. I recall speaking to another Working Group member during that break about a smaller technical matter (communicating signals between servers and end users) and that member expressing skepticism about the value of working out any such details when the real blocking issue was being discussed elsewhere.

Having not been in that smaller side conversation, and not having any notes from it, I find myself both as a participant and as a researcher frustrated and unable to draw an express conclusion of what exactly was discussed. It seems that some anti-trust concern had been expressed regarding an agreement from browser vendors (key implementers of DNT and open to agreement with the DAA proposal framework) to agree to some limits on what blocking they would engage in for online services that complied with user's expressed DNT signals – the concern being that this would inhibit competition between browser vendors on privacy features around cookie-blocking, tracker-blocking, ad-blocking, default settings and how DNT signals were set by users or perhaps restrictions put in place by

¹²⁹Minutes, May 6th

browsers on the installation of extensions that would set DNT signals. That I can't get more specific on what in particular was the concern is a side effect of the lack of transparency and different interpretations from different people I spoke with.

However those details were discussed, Swire and others were not confident that a deal could be announced or agreed upon with this expressed anti-trust concern from the FTC, and the remainder of the Working Group meeting focused on a smaller set of actions the group could take going forward¹³⁰ but without any of the larger deal resolution that had been hoped and planned for.

5.3.2 The different purposes of transparency A lack of transparency about this particular conversation or controversy is frustrating for the researcher, sure, but this example also illustrates some of the different ways that transparency can contribute to the legitimacy of a governance process.¹³¹ Transparency can: 1) establish a record for later debate or review; 2) provide the opportunity to address facts or issues during a deliberation; and, 3) better inform stakeholders about influences or disruptions to a process.

That transparency comes up in the context of anti-trust in a standard-setting body is not unusual; indeed, transparency is a key procedural protection that standard-setting bodies rely on to avoid the liability of potential anti-competitive behavior for their participants. By having meetings, discussions and decisions clearly documented, groups can have a record of their reasoning that can rebut subsequent allegations of anti-competitive motives.

In this sense, transparency is building evidence for later arguments. This is one characteristic of transparency in legislative contexts, where, for example, the Congressional record can build evidence for later judicial interpretation or review. Standard-setting organizations like W3C also use this as a logistical cost-saver: by having records publicly archived, responding to legal threats and steps like discovery becomes trivial – counsel can provide a link to a mailing list rather than exhaustively reviewing private records for relevance. In this case, though, transparency is lacking not around a decision that might have implicated the parties in collusion, but rather around the details of a concern about a decision that was not made.

¹³⁰See the full day's meeting minutes [Minutes, May 8th \(messier\)](#) and the final deliverable agreed on and published at the end of the meeting: [Consensus Action Summary \(deliverable from May 2013 meeting\)](#).

¹³¹This is a very limited subset of the potential uses of transparency for governance generally; consider, for example, the four kinds described by Kosack and Fung (2014).

It seems plausible that the anti-trust objection was misinformed or even simply misunderstood; without having it publicly described, there was no further opportunity by the broader group to analyze or evaluate its significance. Closed-door conversations can be derailed in ways that might have been corrected in more transparent settings. In this way, transparency is a benefit for legitimacy not in uncovering corrupt motives, but instead in allowing issues to be rebutted and responses to be made. This is characteristic of transparency in administrative law, where rule-making procedures typically involve transparency about all data and comments that went into a decision, so that impacted stakeholders have the opportunity to respond.

However, it also seems feasible that the anti-trust objection may not have been the ultimate problem, but rather that ad industry consensus was unstable and couldn't remain around the DAA "Draft Framework" proposal. Under that analysis, the main impact of the anti-trust concern would be procedural or disruptive – it made it harder to get a simultaneous commitment from many stakeholders in May 2013, which subsequently made it harder to implement a DAA and browser agreement, even though anti-trust may not itself have been a large substantive risk.

Taken in that procedural disruption way, it becomes more subjective how to frame the impact: if you want to blame the FTC for the lack of DNT agreement, you can do so; if you'd rather blame ad industry trade associations, you can do that; if you want to blame some other group, you can – there will be little documentation to settle those disagreements ultimately now. There would always be contemporary and retrospective debates about those questions, but it can be qualitatively different when points identified as key by the participants lack transparency. From a retrospective view in conducting research on this process, more transparency might have provided more evidence regarding what would make multistakeholder processes more or less likely to reach consensus outcomes, but that kind of transparency may be different from the transparency necessary to support the legitimacy of a consensus agreement.

5.3.3 Lessons for the policymaker's role What we might be able to take away from a procedural disruption interpretation, though, are some reflections on policymaker participation in multistakeholder negotiation.

FTC representatives deliberately chose to employ a soft touch and communicated with stakeholders more through side conversations and less as leading the way within the process or setting out very particular goals. That appears to be an

intentional strategy to delegate not just technical implementation details but also the political process of finding an acceptable outcome to the stakeholders themselves through the multistakeholder group, rather than being the direct source of the final outcome.¹³² Some participants retrospectively concluded that more aggressive engagement could have moved things forward.

So, trying to get the FTC to be more aggressive in there. I mean, again, it's not the FTC's instinct. I just think the instinct is to lay low, but maybe it shouldn't be.

Many also attribute FTC's use of soft power as driving engagement with the Do Not Track process or with work on DNT at all. That these might be examples of "soft power" doesn't mean they're similarly soft in touch in the sense of being hands off or indirect: Chairman Leibowitz could give quite prominent speeches on the topic.

Whenever the FTC chairman gave a speech about problems in this area or someone senior in the European Commission did, there seemed like there was more interest in dealing

[industry] wanted to know if Leibowitz was going to make good on his threats to take action, and so call it regulation by raised eyebrow or whatever you want to call it.

But that combination of a loud supporter of the process and a quiet on-the-sidelines participant may lead to challenges when the FTC has a concern or discourages a negotiated agreement in process. Positions communicated by policymakers in private fora can have a strong influence, and one that lacks the transparency benefits of administrative law including activating stakeholders and providing a possibility to respond.

¹³²For concerns regarding legitimacy and accountability with strategies of delegation, see "Drawing comparisons" in [Internet Standard-Setting and Multistakeholder Governance](#). This distinction between technocratic and stakeholder-balancing or democratic views also arises in questions over the individual's role in the following section on [Individuals vs organizations in standard-setting process](#).

5.3.4 Different effects on competition Competition also comes up as a theme not just in the specific sense of the collusion targeted by anti-trust law. One perspective that some interviewees emphasize is that market competition is simply a constant background motivation for corporate investment in participating in technical standard-setting.

Standards are a competition, right? Standards are always a deliberate act. [...] People come to the table with vested interest. Everybody at the table in a standards body has an objective that they're working towards, something that they see as an outcome, and there are winners and losers in the standards process.

Regarding the DNT process in particular, interviewees refer to these competition issues, sometimes with aggressive language.

Everyone was using it as a way to get a competitive advantage to screw the other group

I don't always agree with this privacy advocate, but we often can have a discussion, but at least I know where they're coming from. They're not trying to steal my customers or kill my products so they can sell more of their products. But when it's a competitor, you know that they are trying to steal your customers and to kill your product to sell more of their product.

In many cases, the "competitor" is, implicitly, a competitor in a slightly different field (or that entire field or business model), rather than a more direct competing company. So, a third-party advertising network might identify the competitor as the browser vendor whose changes in functionality may affect their business model (rather than another third-party advertising network with a similar business model directly competing for the same customers on a similar basis), or a company with a large first-party presence might be able to siphon advertiser customers from third-parties.

These particular arguments are interesting to me because they seem to put a normative preference for the status quo, a kind of entitlement to current infrastructure. It's 'uncompetitive' in this sense to change the technical infrastructure

that another company's business model currently relies on, but it's not implied that there was some obligation to build the technical infrastructure to be that way in the first place.

Large shifts in technology don't seem to have the same entitlement effect, and so, for example mobile device operating systems don't have to provide all the same tracking features that desktop Web browsers previously had. For example, references to the "post-cookie world" are commonplace in the online advertising trade literature – the terminology can refer to many things, but often includes "mobile" (referring either to mobile operating systems not providing the same cookie functionality or the relative popularity of iOS and Safari which has had stricter cookie limitations), device proliferation or, especially recently, increased cookie blocking from browsers.

Once a system has been around for a while though, making changes leads to calls of anti-competitive practices, although usually between companies that aren't direct competitors. In a way, this becomes a version of backwards compatibility and avoiding deprecation of features: advertisers will regularly lobby Apple or Google to slow down their publicized plans to limit access to IDFA (the iOS identifier for advertising) or third-party cookies.

There is similar sentiment from those who don't identify it as an explicit attack or attempt to undermine others: "it's sort of an interesting shift in power because it would only constrain everybody else." Or, related, that the differences in compliance costs could have disparate impacts based on company size even if a technical and regulatory architecture required consistent application by all competitors:

The cost of implementing things is more easily borne if you are a major corporation. If you are tiny and you suddenly have all of these compliance things to worry about, you can't even get off the ground, right?

These are considerations that are familiar to public policy experts, although in some cases they're being expressed (sometimes with a sense of novelty) by individuals with more engineering-focused backgrounds. Concerns about the impacts of consolidation among tech firms or among firms involved in defining the prominent platforms or protocols for the Internet are widespread beyond this study of privacy in standard-setting. Corporate consolidation and influence in

Internet standard-setting may be described in part through more quantitative analysis of participation patterns.¹³³

But standard-setting is also explicitly identified as a pro-competitive process by some interviewees, and the cooperation between direct competitors is common and notable.

Relationships can become convivial and informal between representatives of directly competing companies:

it's not just they're in their own little world and I'm in my own little world, and the only time we meet on the battlefield is at the W3C, but I'll pick up the phone and call [redacted] and say, "Hey, what are you doing about this?" And it's a very comfortable thing, so I don't have to write an email and proof the email and make sure that it comes from the right point of view.

And that extends beyond the interpersonal level: "it's friendly competition." This comes up in the sense of collaborative development of a platform, e.g. "the Web platform," by competitors – either among browser vendors or of the Web field more broadly. That standard-setting is a competitive act or effects the competitive market between companies doesn't negate that standard-setting for the Internet and the Web is a cooperative act that enables a range of commercial and non-commercial activity. As we will see, that shared perspective among the individuals who conduct that work, but also work for employers competing for customers, helps define the ultimate and often policy-related effects of these processes that bridge diverse and competing organizations.

¹³³See the section on [Who participates and why it matters](#), but also ongoing research work, within the Bigbang project or in the work of Niels ten Oever, including: [ten Oever and Beraldo \(2018\)](#) and [Arkko et al. \(2019\)](#).

5.4 Individuals vs organizations in standard-setting process

5.4.1 “the theory” of individual participation

Individuals who participate in the process are the fundamental unit of the IETF organization and the IETF’s work. The IETF has found that the process works best when focused around people, rather than around organizations, companies, governments or interest groups. That is not to say that these other entities are uninteresting - but they are not what constitutes the IETF.

– A Mission Statement for the IETF ([Alvestrand 2004](#))

The procedural principle of individual affiliation is frequently cited and discussed by participants in IETF standard-setting, as well as in its documentation and by its leadership. Individual affiliation takes on a sort of mythical status: everyone knows and talks about it, and knows that it isn’t quite true, but also that it has some weight and history behind it. According to my interviewees with substantial experience at IETF, it’s a “narrative” or a “story” or a “theory”:

there is the theory that all people at the IETF are participating in their individual capacity and are not representing their employer. I say that is a theory because of course in reality most people there are acting consistently with the interests of their employer, but, I mean, especially in the early days of the IETF kind of back into the '80s and '90s [...] engineers who were there [...] their employers were forward-looking enough to assign them to essentially go contribute to the IETF and to not be heavily pursuing a corporate agenda.

Participants note examples (sometimes very specific, sometimes general) where an employee will take a different position from their employer’s direct interest, both at IETF and W3C, or where they themselves handle the tension of aligned but not identical goals for their work. These separations are attributed to a few factors:

- “conscience” from an individual doing the right thing counter to their employer’s interest,
- autonomy and flexibility that an employee may be granted perhaps related to seniority,

- credibility of reputation developed by individuals who participate while employed by multiple organizations over time

Sometimes you see someone who does something that's clearly not in their employer's interest and you just think, "wow, that's amazing." And sometimes there are really serious experts who are distinguished enough that they have a lot of freedom in their job, or they work in research or a lab or something and they're not constricted by the fact that they already have features built into a product.

That was [Company's] goal. Yeah. It's not my goal. My goal was to not break the Internet. It's most of what I do, is not break the Internet.

Employees at times have this latitude because of an interest from their employer in supporting standards development work: sometimes that's because the company has a particular product or business goal dependent on improving standards in an area, but also it can be because a company wants "insight into what was going on" (or similarly "active awareness") or "visibility" by having their employee in a prominent role with "active involvement."¹³⁴ There are significant similarities here to the boundary organization collaboration of open source software foundations (O'Mahony and Bechky 2008): where companies can fund employees to work on "areas of convergent interest" in an open source project, while allowing that project to maintain its own practices.¹³⁵

5.4.1.1 Effects on behavior The examples above of individual participation share a sense that an individual's role in the larger Internet community has, sometimes, a priority over their role as an employee or representative. Long-time standard setting participants note that this experience can influence or moderate how an individual behaves because of their longer-term, cross-company interest.

Most people don't stick to a single company forever. And so, that colors behavior. That works as long as you are within what is typically called a community or within an industry in the broadest sense.

¹³⁴These quotes in passing are from two different interviewees.

¹³⁵This connection has been detailed previously in Doty and Mulligan (2013), and there are particular connections here since there is overlap between those open source projects and standard-setting organizations.

Engagement in the collaborative process of developing the infrastructure that is the Internet requires or at least benefits from this long-term and larger-scale commitment. These situations in technical standard-setting seem especially relevant for engineering professionals, rather than others in business, product development or legal teams, where financial considerations or the client's position are more primary to the job. And it coincides with technical standard-setting organizations being an arena where engineers are often given latitude to take public positions without requiring more extensive sign-off from the rest of their organization.

The closest analog that I heard outside of engineering was from ad industry participants who sometimes refer to the future direction of the industry as a whole. In particular, this seemed to be attributed to leadership, in the sense of senior executives of large companies or leadership of industry trade associations, rather than a perspective of typical participants or employees of member companies.

And so you can cut through sort of, you know, junior-level perceptions of what they feel their goal is, and by moving it to that board level, they truly do care about the ecosystem, right? They want all businesses to flourish.

A focus on ecosystem or platform or more than a single company or interest can be compatible with this sense of forward-looking planning and more collaborative behavior. In the case of high-level firm leadership or technical seniority, it might in both cases be motivated by the possibility of moving between firms and organizations and having an interest in positive relationships and reputation. Consider the trends described in *Regional Advantage* (Saxenian 1996) and the Silicon Valley culture of employees easily moving, and cross-pollinating ideas, between competing chipmakers.

However, a more cynical explanation for an individual's divergence from an organization's interest comes up in the context of Do Not Track:

there's a class of kind of advocates, mostly on the industry side, probably entirely on the industry side, at least from my point of view, who have a vested interest in gumming things up and in friction when it comes to efficient government action, and, honestly, a lot of times they're not doing it, I think, even in the best interests of their clients or particular companies. They're doing it because it makes them more important. It raises their profile, makes them more essential. It makes them the kind of main roadblock, which makes them the main broker

that you have to pay attention to, and of course it leads into billable hours. And so I feel like there's this class of D.C. advocate, D.C. lobbyist who fits that bill, and I think I felt like the Do Not Track room was full of enough of those people that it derailed a lot. I know I'm making a lot of assumptions here that probably other people would disagree with, but that was kind of my take.

That is, while we might identify cases where an individual will diverge from an employer's direct interests because of a community interest, there may also be cases where an individual may diverge in order to benefit themselves directly, in terms of power or financial interest. In that case, divergences can lead to more obstructive behavior, relative to what the client might prefer, or relative to what the broader community may need.

5.4.1.2 Individuals representing an organization And in the context of negotiations over Do Not Track, there are also objections to these divergences or a preference or respect for direct representation of an employer's or client's interests. Respect for representing a client or a position is tied to the idea of being a good faith and principled participant who believed what they were saying or advocated for a genuine interest, even when the speaker disagrees with the identified person.

two or three people in the Do Not Track process, who were on the other side of me on every issue, who were not terribly pleasant to work with, but at the same time I think were basing their opinions strongly on ideals or strongly on what they thought one particular company wanted or needed

We almost never agreed substantively on the issues, and [he] might skewer me about half a dozen things in the press or on Capitol Hill. [He] and I have testified on the same panels in front of Congress. [...] We had a job to do. We both believed in what we were doing, and I think we both did a really good job.

These are positive evaluations of behavior not just because someone's statement might be reasonable or principled, but because they were advocating for a company's or an organization's interest and that was their "job to do."

(See also: good faith vs bad faith and participant antagonism, as detailed in the [earlier section on process](#).)

Under the perspective that employees should represent their employers faithfully, there are also cases where a lack of internal coordination can create a situation where a representative doesn't know what other decisions may have been made internal to a firm. This can cause a change in position in a standard-setting negotiation or even conflicting statements from individuals in different teams who are employed by the same company. While this doesn't get described in the same way as bad faith, it can lead to practical frustration.

People in the room have views different than their companies. [...] the people who were representing Microsoft at DNT did not know that Microsoft was about to change its position [re: default settings].

Or, describing inconsistent positions and a lack of shared information from experience (this quote refers to employees from a different tech company):

he was pretty down on the lack of specificity of the API and I was kind of like, you work with the person who edited this thing, right, really? [...] honestly, what are you doing here? Can you guys go in the corner and strategize for a little bit before you get up and say this? [...] Your coworker over there who's actually building the product thinks the exact opposite of you.

This phenomenon is common in standard-setting, though it's ironic to hear Microsoft as one example, as the informal reputation I had seen built up was that Microsoft was more likely to discuss and agree on something internally before taking a position at a standards body, where Google was more likely to have different teams that didn't coordinate about their products and employees might take opposite positions at the standards body and resolve them after the fact. People have preferences in either direction – where the lack of internal coordination can be frustrating, it can also be appraised for speed or transparency – but organizational cultures differ on how much coordination happens internally and separately from the more public standard-setting process.

Some have positioned themselves as fundamentally opposed to companies that do not sufficiently coordinate employee actions. Regarding potential changes to cookie functionality in Firefox, the Interactive Advertising Bureau CEO Randall Rothenberg described Mozilla in a trade press interview (Ebbert 2013):¹³⁶

¹³⁶This quote is *not* from a research interview and is not similar in style to a research interview discussion; incendiary statements of this kind are part of active campaigning, and not uncommon from Rothenberg.

Mozilla is obviously a very factionalized organization. It's like mob rule. It's very difficult if you're a rational player. [...] It's not really clear if there is a Mozilla itself, other than the radical players who seem to have the ability to control what does or does not go into the browser.

This “mob rule” reference is perhaps frustration regarding the model of open source software development or the ability of non-employee developers (here called “radical players”) to contribute code to the Firefox browser. Perhaps analogous is similar frustration expressed (again, with caveats, in the press (DePillis 2013), from an advertising industry lawyer) about the input of developers and developer associations in a NTIA-convened multistakeholder process about mobile app transparency: “Developers are the people we hire to do software coding. That's like saying the painter of a retail store makes decisions about the paint.”

That engineers or developers may contribute to decision-making about software, and may have the ability to do so without the hierarchical approval more common in other industries, may contribute to some of the surprise or mismatch of expectations here. As described in [The Ethics of Engineering](#), previously, engineering is inherently ethically-laden and there is an impulse towards integration of larger concerns into the practice of engineering. If engineering is inherently about deciding on and bringing into being the good life, then the choices of engineering will be significant in a way that isn't simply the direction of their employer.

5.4.1.3 Challenges to the equality of individuals In addition to a concern over individuals not acting in concert with their organizations, some participants seem to consider the equivalent status of individuals as an affront, or impractical, because of its informality or the lack of appropriate representation.

And so you had somebody there representing a 100-billion-dollar industry and you've got Jonathan Mayer who was what, 22 at the time? And very smart. But the idea that he was going to show you, it was just ... it was really – there were so many dynamic issues at play.

It can at times, as in this case, be difficult for participants to put directly into words this mismatch feeling. It may also be related to discontent (this same

participant calls it “absurd”) with the informality of procedural steps like a “hum” to gather a sense of the room.¹³⁷

Lower barriers to participation may be preferable for the idea of access, a potential boon for legitimacy. But low barriers may also prove frustrating because of who may choose to participate.

one thing that’s very odd, is who’s on the call and who’s not on the call, right? So, if an individual in some country decides to show up every week, they get to speak. And there’s a sort of formal equality to their participation with the participation of people who’ve invested a billion dollars in a particular thing. Or, the same as one of the public interest groups that’s invested years and years of effort in the space, and then some loud-mouth, who just like feels like showing up, gets the same formal role in the process as the dedicated group that’s well-staffed and very thoughtful. So, one way to say it is, you don’t have to have skin in the game to be on those calls

“skin in the game” is an especially evocative description of this challenge, as it provides a similar metaphor to “stakeholder,” that some commitment, investment or ownership is a reason itself for legitimacy in participation. Individuals may not have the same ability to demonstrate that weight that someone affiliated with an organization can. This is a sharp contrast to the legitimacy concerns raised about implementers (and especially market leaders) having too great a power in an interoperability-focused consensus process.

5.4.2 Different views of individuals in multistakeholder processes: representational vs collaborative One view of governance is that of balancing the interests of stakeholders: what makes a decision legitimate and valuable is the positions of important sub-groups that have a significant stake in the outcome. Distinct, and often quite different, is a process where legitimacy and value comes from efficient and effective analysis of arguments’ validity by the key experts in the area – we might call this the technocratic view.¹³⁸ Internet standard-setting has typically taken a part of both approaches through its interoperability focus: implementer

¹³⁷Mentioned in brief in a popular press piece about the Do Not Track process: “Do Not Track’ Web System Stuck In Limbo” (2012).

¹³⁸Sunstein describes the tension, which applies to all public law but especially administrative law, as between the technocratic and democratic view, of expertise and accountability to the public will (2014).

weight is especially important in determining acceptable outcomes but insight is valued for finding and evaluating technical feasibility.¹³⁹

That individuals participate in a process rather than organizations is on one level obvious: it's people who are in the room or on the calls who are talking and debating. Companies and organizations don't take actions like that, although it's very common to anthropomorphize them in our language, perhaps based on a mental model of firm hierarchical decision-making.¹⁴⁰ That individuals are affected by their affiliation with organizations, influenced by corporate priorities even if distinct in their own goals and how they interpret them is also clear. What's of interest, though, is how those individuals within (or perhaps mediated by) organizations build relationships, collaborate, compete and come to decisions in a consensus standard-setting process.

The representational view of individuals in a multistakeholder process minimizes their personal perspectives, expertise or interests: the role of a representative is to accurately represent the positions she has received from the represented group and to argue most effectively for the represented group's interests. And for the stakeholder-balancing view of multistakeholder process and its legitimacy, representation is a natural fit for what's expected from each participant.

But a collaborative view of individuals in a multistakeholder process maximizes their expertise in a problem-solving orientation, where they have a high-level goal based on (or at least aligned with) their organization's priorities and have autonomy to use their own perspective to navigate towards a practical outcome. For a technocratic view of processes where quality of argument is key, encouraging individuals to collaborate with their own expertise seems most fitting.

I argue that the Internet standard-setting process attempts to accommodate both the stakeholder-balancing and the technocratic view of the process *and* both

¹³⁹That it might be possible to bridge these procedural and substantive views of legitimacy, or both rationality of argument and the material conditions of the world in making decisions, is not novel. For example, Habermas' view of discourse ethics maintains that under the right conditions of autonomy and speech, procedural requirements can guarantee just outcomes, and Froomkin argues that IETF standard-setting is a rare (unique!) example of this idealized practical discourse within a particular community (2003). This work is not focused on testing this hypothesis, although there is plenty of evidence here about strategic action to counter that argument if others are interested in it. Instead, I am simply trying to explain the different approaches that people have to a multistakeholder process and the sources of the conflicts for those who do not accept a combined view and how they respond to individuals and their unclear relationships to larger stakeholder organizations.

¹⁴⁰See also the concept of "institutional synecdoche," as described previously in [Internet Standard-Setting and Multistakeholder Governance](#).

the representational and collaborative views of participation. There is audible frustration from anyone expecting purity of either perspective. A representationalist will be angry that an individual without sufficient stake needs to be addressed; a collaborativist will be disappointed that crass concerns about someone's business model interfere with a more rational or beneficial outcome. But the potential benefits of using consensus multistakeholder technical standard-setting to make progress on tech policy challenges rely on both the pragmatic feasibility of finding outcomes acceptable to key stakeholders and the potential innovation of problem solving among a heterogeneous group of people with varied backgrounds and expertise.

It is possible that to the extent that the architecture and implementation of the Internet has been a success for liberal – in the sense of Postel, but also perhaps Mill – outcomes, it is the combination of those views of a process for individuals that has contributed to its success.

5.4.2.1 Individuals vs organizations There is an imperfect but apparently substantial alignment between organization-centered and individual-centered perspectives on process. Consider the following categories.

Table 2: Individual vs organizational views of multistakeholder process

	organization-centered	individual-centered
unit of participation	organization	individual
scope of work	policy	technology
task	decision-making	implementation
purpose of a process	balancing of stakeholder interests	technocratic
form of process	negotiation	problem-solving
role for individual	representation	collaboration

Not all disagreements or conflicts described in the qualitative analysis above line up with these categories. For example, some might agree with the individual's role as representational, but disagree about who should be represented (herself, one's employer, the larger cross-organizational community to which one belongs). But the merging of these two categories does explain some regular confusion. Some object to a technical standard-setting body as the wrong venue for making policy decisions, preferring legislative or administrative governments for that purpose; others object to the engineers being involved in decision-making rather

than sticking to implementation.¹⁴¹ The tension of impulses towards separating engineering from ethics or policy and the impulse towards more directly integrating it reflects this distinction in the particular field of engineering.¹⁴²

One approach is to explicitly recognize the difference between these categories and that they're both relevant and then try to divide them up. P3P designers tried to make that explicit with vocabulary vs policy choices, and the Tracking Protection Working Group also divided deliverables into compliance and preference expression syntax. Some involved in the DNT process and some that I spoke with would have advocated to push that separation further, separating technology and policy and enabling a tussle over which compliance policy would be chosen or accepted.

What if, instead, we recognized not only the differences but also the increasing blurring of these boundaries and embraced that merging? We might, as in the case of DNT, see processes where stakeholders discussed and learned technical details, business practices and policy implications in a combined setting and we might see more people with combined technical and policy expertise, as discussed in the following section on patterns of participation.

¹⁴¹There are also contrasting views of engineering as a profession and the relationship of the individual engineer to her employer. Following that representational view, the engineer works on behalf of the client, implementing to their exact specifications and setting aside any of her own value judgments for the needs of the client (which might be the employer, or the end user). Following the collaborative view, the engineer considers not only the needs of the client, but also her own values and her insights into the design and how it will interact with others.

¹⁴²Again, see, previously, [The Ethics of Engineering](#).

5.5 Who participates and why it matters

In studying processes like technical standard-setting, I have been especially attuned to who is participating. In order to evaluate multistakeholder processes for developing techno-policy standards that can resolve public policy disputes, we must consider access and meaningful participation – essential criteria for both the legitimacy and the long-term success of these governance efforts.

But who participates will be measured not just by personal characteristics, but also by the political importance of the stakeholders who are represented in a particular process. How the stakeholder groups are divided up and the number of “sides” they represent is discussed below.

Participation is not a binary, in-or-out characteristic, so this section also looks at the roles that participants play within technical standard-setting communities: who stands out as formal or informal leaders and how the social network is structured.

Finally, I look at the expertise and experience of participants in developing techno-policy standards and how participants call out the need for more integrated backgrounds.

5.5.1 Why participation matters for legitimacy Informal and non-state processes may have opportunities for more open doors and greater access by anyone interested or affected – you don’t have to be elected or pay a large fee to show up at a conference call or mailing list – but they also may (and often have) not. In addition, voluntary standards aimed at interoperability have a certain kind of legitimacy backstop: if the implementers aren’t in the room (a failure at the step of convening¹⁴³), then it’s likely none of the durable effects of a standard will be implemented.

But the scale of those affected by the future design of the Internet is extremely broad and not limited to the companies likely to implement any new standard. As described previously, a consensus for interoperability may be meaningful, but alone won’t settle concerns about legitimacy.¹⁴⁴

The diversity of stakeholders for the Internet and Web is enormous – including governments and businesses of all kinds, as well as end users from around the world. Who participates and the industry or organizations they represent may

¹⁴³See the [stages of success in standard-setting process](#).

¹⁴⁴See [Internet Standard-Setting and Multistakeholder Governance](#) in the section “Legitimacy and interoperability” and in [Doty and Mulligan \(2013\)](#).

determine how technology is designed and what functionality and values the larger socio-technical system provides. And for questions of direct public policy importance, the legitimacy derived from participation may have a greater weight than it is on matters that appear to be more simply technical or functional.

5.5.2 Demographic representation in technical standard-setting That the experts participating in the detailed technical standard-setting processes – including ones specific to key issues of online privacy – are not representative of the world, or the United States or the users of the Web is well-known and widely accepted. Feng, for example, asks “where are the users?” and argues that serious limitations arise from end users not being able to effectively participate and not necessarily being either well-understood or well-represented (2006). Froomkin (2003), even in arguing that IETF practice is a form of ideal discourse, raises the question of, “where are the women?”¹⁴⁵ Froomkin accepts that the IETF is dominated by English-speaking men, but hints that diversity may be improving because of a woman in a position of leadership; no quantification is present.

One limitation noted by these two particular authors but encountered whenever the problem is raised is that effective participation in standard-setting fora covering these detailed technical topics requires extensive expertise, as well as time and money. As noted previously,¹⁴⁶ while formal barriers may not prohibit anyone from reading mailing lists or joining teleconferences and while fees may not generally be prohibitive, the time involved to read every email message, the money to spend those hours and to travel to in-person meetings in order to be most effective and best connected to all other participants and the training necessary to understand the implications of proposals or to recommend alternatives are all limiting to general participation.

However, some participants in the Do Not Track process also noted that expertise regarding Internet architecture would not be the only or appropriate kind, in part just because of the lack of demographic diversity. That missing expertise might include not only particular disciplinary training in ethical, legal or policy issues but also cultural or personal understanding of lived experience.

I don't even know how to frame that debate [over what is ethically acceptable re: privacy], and I think having technologists try to work

¹⁴⁵In this case, quoting and citing feminist scholars who are critical of Habermasian discourse theory.

¹⁴⁶See Doty and Mulligan (2013) and [Internet Standard-Setting and Multistakeholder Governance](#).

out the answer to that kind of question is horrible. We need ethicists and lawyers and sociologists and so on, people who understand social debate and policy and norms to have that debate. I also think that technologists having that discussion will be culturally insensitive; the bulk of the technologists are Anglo males, perhaps not the bulk of the world's population are affected by this debate.

This perspective may seem familiar; there is an argument that the profession of engineering may rely on a higher ranking of “poets, philosophers, politicians” to settle fundamental questions of values and that there is a separation of concerns between engineering and analysis of ethical values.¹⁴⁷ Our correspondent, a technologist in their own framing, distinguishes that issue of policy expertise from cultural sensitivity and demographic representativeness, but the ideas are intertwined.

This was echoed by another participant who tied the specific lack of gender diversity in meetings to a concern that Do Not Track or related privacy work involves policy goals, despite being a technical standard. While there are significant reasons to be concerned about the lack of gender and demographic diversity in engineering communities in general,¹⁴⁸ diversity of participation is identified as especially important for questions of policy or ethical values.

5.5.2.1 Semi-automated estimates of gender and participation There are many demographic dimensions that may be relevant to questions of legitimacy over the design of Internet protocols. Because these tech communities face prominent controversies over sexual harassment and discrimination in employment contexts, gender has been one such area of interest. Gender is: 1) highlighted by some interviewees as an important demographic characteristic with a marked disparity, and, 2) an area where we may be able to use quantitative data to validate and

¹⁴⁷Ortega y Gasset and Miller (1962), as detailed previously in “Separation vs. Integration” in [The Ethics of Engineering](#).

¹⁴⁸Namely, at least, the following:

1. Equality of access to opportunity in engineering careers
2. Quality of heterogeneous teams
3. Relevance of personal experience and knowledge in engineering
4. Inherent and essentially ethically laden nature of all engineering (as described previously in [The Ethics of Engineering](#).)

explore the disparity at a different scale. As such, it's a fitting particular case to explore with a mix of methods.

5.5.2.1.1 Methods, questions and caveats Mailing list conversation represents a primary discussion forum for IETF and W3C standard-setting conversations, including the Tracking Protection Working Group, and these mailing lists are publicly and permanently archived. Using those mailing list archives, we may begin to gather data on questions such as:

1. What is the gender distribution of participants in Internet and Web technical standard-setting?
2. How do gender distributions vary between different groups?
3. And, in terms of evaluating the practicality of this methodology: to what extent can fully automated or semi-automated methods be used to provide estimates of gender distribution on large, computer-mediated communications fora?

These are relevant and important questions for the larger project's attempt to understand patterns of participation and what conclusions we can draw about representation and legitimacy of decision-making. For the utility of metrics for demographic diversity in large data sets, the caveats and ethical considerations in conducting that analysis and in the automated methods for doing so, I have tried to build on the work of J. Nathan Matias (2014).

Like all methods, there are substantial limitations in using quantitative, automated tools. Significant caveats must accompany the use of these tools for measuring the demographics of participation.

- Identifying individuals in computer-mediated fora is difficult. There are few restrictions on the names or email addresses that participants use, people may use multiple email addresses at once or change them over time or share them.
- Inferring gender, through automated or manual means, is known to be imperfect. Neither automatic inferences nor human annotation will always accurately identify someone's gender.
- Gender is neither perfectly stable nor ultimately externally observable. The presented gender of a participant may change over time and may not be known by other participants or an outside observer.

- Cues for gender vary across cultures. While names, pronouns or other language use may be specifically gendered in some languages or nations of origin, that may not apply in all cultures.

These caveats provide context for the interpretation of results. In particular, this method is not a reliable way of determining a particular person's gender. While intermediate data files will include an inferred gender for many people, individual values are not presented in results and should not typically be used. In addition, population-level results may be skewed based on how people choose to present themselves in these online technical discussion fora or based on limitations of either automated or manual methodology.

If the caveats are so significant, is this work still worth doing? I believe that it is, for these purposes:

- descriptively evaluating the demographics and representation of decision-making groups where participation is considered important for legitimacy;
- generating trends or identifying anomalies that would benefit from further investigation; and,
- evaluating the utility of automated and semi-automated methods for estimating gender and other demographic characteristics in computer-mediated fora such as mailing lists.

To estimate the proportion of gender of participants on standard-setting mailing lists, this work uses BigBang¹⁴⁹ to crawl, parse and consolidate mailing list archives. The automated analysis here makes use of Gender Detector,¹⁵⁰ a library which makes estimates based on historical birth records, as described by Matias (2014). Gender Detector is configured to return an estimate only when those birth records show a very high correlation that a person with that name is assigned that gender.

5.5.2.1.2 Initial results on gender disproportion Further analysis of semi-automated methods and different levels of manual resolution will be addressed in future work. But for this initial investigation, we can review initial results from the automated process, for some insight into the three questions above.

¹⁴⁹<https://github.com/dataactive/bigbang> and see the [Methods chapter](#) for more description.

¹⁵⁰<https://github.com/malev/gender-detector>

Of the dozen groups with the largest fraction of participation from women (these are the above average groups in this dataset), the dominant topics are: accessibility, publishing and privacy.¹⁵² Even prior to any further manual annotation, this data suggests that standard-setting around Web accessibility in particular may be less male-dominated than other Web standards topics. That the privacy topics (the TPWG and PING, the Privacy Interest Group) also see relatively higher participation from women in initial analysis might be a prompt to explore whether the gender diversity is more significant in an especially policy-relevant area. Further qualitative work to investigate this demographic difference in particular groups could be rewarding.

Finally, in every group there is a significant fraction of the participants where we can't automatically estimate the gender. In particular, groups that discuss internationalization or for other reasons have higher participation from Asia are especially difficult to estimate, as the automated system is configured based on US-based birth records. This limitation was known prior to any data analysis, but it's notable that because different groups may have substantially different makeup by countries of origin, estimating gender based on name may be difficult to compare. If further automated or semi-automated methods like this will be used, it might be worth exploring combining datasets that could detect gender equally across multiple countries of origin rather than assuming that Western data will be dominant.

5.5.3 Stakeholder groups: counting sides Another way to answer the question of who is participating in a standard-setting process is to identify the particular stakeholder groups they represent. As I sketched out previously,¹⁵³ stakeholder groups in standard-setting bodies overlap, and individuals can be affiliated with different sectors over time. People I spoke with who participated in the Tracking Protection Working Group or were involved with Do Not Track more generally were sampled from these different stakeholder groups, but I also prompted them to discuss the other stakeholders they were interacting or negotiating with.

Of particular interest, was the idea that there were two sides in the debate over Do Not Track, a theme that arose during conversation with several participants despite it not being one of my prompts.

¹⁵² Also in this selection: a group that discusses life science research, a group that discusses technology for “second screens,” and a group that discusses XML stylesheets.

¹⁵³ *A Mixed-Methods Study of Internet Standard-Setting*, in particular “The networked site” and “Dimensions for sampling” although the ideas come up throughout the methodological overview.

The biggest problem with DNT is it was set up as trying to find a compromise between [...] privacy researchers and privacy advocates on one side, and advertisers on the other. Privacy researchers have no incentive to let companies gather information about their customers. None at all. No reason for them to. The advertising industry has no incentive to take care of the customer or reduce the amount of data that it collects. No incentive at all.

This participant identifies the two-sides framing as more extreme participants, with, as a result, “no incentive” to compromise. Even among more positive assessments, there is a similar view of sides: “I think on the positive side there’s been a tremendous amount of progress made just from a high level in terms of getting both sides to talk.”

However, what the two sides consisted of was not always consistent. Consider two narratives. In the first, Do Not Track is a struggle between privacy advocates and the online advertising industry. Advocates want to promote a new consumer choice tool (or, based on your perspective, want to undermine or destroy the business practices of online behavioral advertising or market research) and compel advertising services to respect it; the ad industry wants to protect existing business models and the economic benefits of ad-supported online content. Obviously any brief description like this is going to oversimplify, but notably this doesn’t mention web browsers (who build the software that sends DNT headers) or web sites (who operate servers that receive DNT headers, and who sell advertising). In the second narrative, browser vendors are building and promoting DNT as a privacy feature for their users (or, depending on your perspective, an anti-competitive move to prioritize their business models over targeted advertising), in opposition to the online advertising industry (that funds much of the revenue of the browser vendor firms). In this telling, consumer advocates are sidelined, policymakers are unimportant and web sites remain uninvolved.

While the former perspective is probably more commonly ascribed in my interviews, the latter perspective is also significant, and gives a very different tenor to the negotiation.

we allowed the debate to polarize like that which I think was not helpful, you know it ended up indeed often with the browser vendors on one side of the table and the ad industry on the other, and the consumer advocates being ignored.

Some explicitly chose to identify a more diverse set of stakeholder groupings as an attempt to unblock negotiations. Peter Swire, in particular, describes five “camps”: “the privacy groups, the browsers, the first parties, the third parties and the regulators.”¹⁵⁴ Some of those terms of art may be opaque: “first parties” refers to web sites, web publishers, online platforms – the New York Times, or Wikipedia, or, often, Facebook are prominent first parties – organizations who operate web services that a user will directly visit; “third parties” – an analytics service, the online ad network that chooses the ad to show beneath a blog post, or Facebook when it shows up as a like button on an article – are embedded observers of such a visit, who collect data about a user’s visit and insert relevant advertising or other content into a web page.

Notably, viewpoints of two sides also come up from multistakeholder process participants I spoke with who weren’t involved with Do Not Track at all, for example: “the business side or [...] the privacy side,” or distinguishing between implementers (especially browser vendors) and user advocates in a W3C context.

5.5.3.1 Relevance of stakeholder group analysis Is the level of granularity really so important?¹⁵⁵ A two-sides perspective can influence:

1. the practical effects of attempting to find consensus;
2. our retrospective understanding of the different viewpoints and dynamics of participants; and,
3. future attempts to design similar privacy controls.

Regarding the process itself, a two-sides perspective encourages entrenchment of participants and seeing the process as contentious. Participants describe a “polarized” environment, and a lack of incentive to compromise or disagree with others in one’s “side” even where there were significant disagreements within industry or advocacy, say. (For more, see [findings on process](#), regarding animosity and agreement.)

Regarding research on the Do Not Track process, a more granular description of stakeholder groups is important for purposive sampling.¹⁵⁶ For example, assum-

¹⁵⁴This is attributed because it is representative of his public approach to chairing and other interviewees also recognize his identification of a larger number of groups, although they might refer to three rather than five.

¹⁵⁵It’s possible to carve up any large group into different numbers and sizes of subgroups, with no essential preference beyond what’s pragmatic for analysis (Quine 1951).

¹⁵⁶Again, see [methods chapter](#).

ing industry (or even, ad industry) as a singular group would have given me a very different set of perspectives if I had only interviewed advertising trade association participants or only browser vendor employees. Beyond sampling, it's useful for research both to recognize variations within these larger categories and also the tendency to agglomerate into two-sides perspective.

Regarding future designs, authors of the Global Privacy Control (which follows a very similar design to the Do Not Track HTTP header) identify first-party publishers as the recipients of the user's expressed preference (one spec editor is a representative of the New York Times) and more explicitly ties the design to specific state legislation. Whether that effort is more likely to be successfully adopted isn't yet clear, but the differences rely on the debate not being as simple as industry-vs-advocacy. Recognizing the multi-party nature and the relative subtleties may help organizers of future multistakeholder process identify distinct and promising opportunities for cooperative effort.

That two-sides narratives also arise between implementers and privacy-advocate non-implementers provides a cautionary tale about the efficacy or legitimacy of these multistakeholder processes. If privacy advocates cannot identify allies among implementers of technical designs, then technical standard-setting processes or other multistakeholder processes where technology is the primary implementation are likely to be disappointing. If organizers of multistakeholder process want the potential legitimacy that comes from consensus standard-setting, expanding beyond reluctant implementers and non-implementing advocates may provide better results.

5.5.4 Roles within communities To understand participation, we have to see not just who is and isn't present, but something of the roles and connections they have within the technical standard-setting process.

5.5.4.1 Leadership One "founding belief" of the IETF, for example, is the lack of formal governance structure: "we reject kings, presidents and voting; we believe in rough consensus and running code."¹⁵⁷ While kings and presidents may not be present, people I spoke with consistently highlighted the importance of leaders, formal or informal, in directing work and ensuring key values.

¹⁵⁷This is prominently described in the Tao of IETF ("[The Tao of IETF: A Novice's Guide to the Internet Engineering Task Force](#)" 2018), quoting David Clark.

For example, some identify the seniority of Area Directors and the process of IESG approval as essential to security and privacy considerations in Internet standards:

the security area directors are like a force to be reckoned with at this point.

IETF leadership have also used the ability to put conditions on the creation of new groups to make sure privacy is considered early on (rather than just at the stage of approving the final output).

the leadership of the IETF in a somewhat unusual move said, “no, you cannot charter a working group to address location unless you address privacy”

A few of the people I spoke with specifically cited the geopriv working group that directly considered privacy, as well as formats for communicating location data. Geolocation is also cited as a key privacy-related datatype in part because of the relatively early development of the technology in Internet and Web standards.¹⁵⁸

Leadership is also often referred to in chairing any particular group, whether a Working Group at IETF or W3C, or multistakeholder processes in other settings. While some participants with experience in such roles describe a necessity for neutrality about both the participants and the outcome, some also explicitly balance that with needing a particular direction or motivation to be pushed forward. This description was given in the particular context of an ad industry trade association process, but applied more broadly, and there are similar phrasings from other multistakeholder process participants I spoke with:

you must have a strong leader with a vision, a goal and an agenda to make any kind of multi-stakeholder process work. In the absence of that it's not gonna have an outcome that I would suggest is beneficial. People may or may not disagree, but I have never seen a sort of multi-stakeholder kumbaya thing produce something without a very, very strong vision and leader who said “This is where we want to get to and try and get there,” understanding you may not get everything you want but set an agenda.

¹⁵⁸This is discussed further in [How participants see privacy](#).

Statements from quasi-leadership organizations and prominent individual contributors have been significant in responding to Snowden revelations about the exploitation of security vulnerabilities in Internet and Web standards.¹⁵⁹

5.5.4.2 Standard-setting organizations as social networks Because IETF is a long-running effort and involves many distinct but connected areas of work, conversation and debate, it's also possible to identify the roles of individuals and the connections between groups as a social network.

To continue with the idea of leadership, a bipartite graph of the participants and the different working groups at IETF makes it possible to calculate measurements like centrality (L. C. Freeman 1978). The people with the highest closeness centrality are the ones that have the most co-affiliation with every other person, or the shortest path to every other person. Automated accounts are, as we might expect, extremely high on this measure – they're used to send announcements of publications and do so to basically every group. The individual people highest ranked on this measure include Stephen Farrell, Jari Arkko, Ben Campbell, long-time participants with leadership roles. The highest ranked woman is Alissa Cooper, current Chair of the IETF.¹⁶⁰

This graph of working group mailing lists and frequent senders can also demonstrate the structure and interconnectedness of these groups, based on the participants who bridge them.

Further work is needed to quantify the relative level of interconnection (what is the appropriate null hypothesis to contrast with?), but the visualization shows that *most* groups and most participants are tied together by these overlapping participants, with just a few more isolated individuals who frequently participate but only on a single topic. In many cases, a multistakeholder process convened to address a new topic or new idea may not have that consistent, multi-venue interaction. Many of the people I spoke with who participated in the Tracking Protection Working Group and the standardization of Do Not Track described it as their first experience with standard-setting or W3C at all, and longer time participants described that as an unusual experience:

¹⁵⁹This is described in more detail in previous work, looking at the responses to Snowden revelations to illustrate how the IETF reacts to exogenous events and how that's visible in mailing list traffic patterns and published documents: Doty (2015a).

¹⁶⁰Annotated code available in this notebook, which also evaluates some other participation metrics for IETF: <https://github.com/npdoty/bigbang/blob/ietf-participation/ietf-participation/IETF%20Participants.ipynb>




Figure 16: Colored bipartite graphs of mailing lists for IETF Working Groups and frequent senders.

So that was a little strange. We spent a lot of time talking about what we could and could not do as a working group, which you don't usually have to do.

Uncertainty about the process and how it works is one area (as this participant is describing), and professionalism and empathy in relationships is another (see previous section on [effects on behavior](#)). But the ability to build ongoing working relationships has also been identified as an important success criterion and condition for long-term success of the coordination we would hope to find from techno-policy standards.¹⁶¹

5.5.5 Expertise and experience A final way to answer the question of who is participating is to describe the expertise and experience of participants – that is, not just who you are in some sense, or who you represent in some sense, but what you know or how you work.

Described as of particular value are those individuals with both technical and policy expertise: because of the tightly intertwined technical details and policy implications of what we have described as techno-policy standards. We have recognized some prominent participation in Do Not Track as having technical and policy experience and described the growth of a community of practice with interdisciplinary expertise around privacy ([Doty and Mulligan 2013](#)). Where individuals don't have that cross-disciplinary background, it might take some close teamwork.

I think that when I participated in P3P, I had an engineer sitting with me. I hired away [...] from a product team a guy who, you know, helped me and I helped him and we were hopefully effective together. But it often requires people with a combination of both, you know, law, policy and technical chops and there's not a lot of people who have both those [...] so it may require a team, you know, unless you're a kind of a standards person who has kind of got a mix of those things.

It was not uncommon¹⁶² for a participant that I spoke with to describe their previous experience with engineering or technology despite working as a lawyer, or vice versa. That additional expertise was often considered a competitive advantage or a way to have more effective input on the discussion.

¹⁶¹Emerson et al. (2009), as described in [Doty and Mulligan \(2013\)](#).

¹⁶²Indeed, it sometimes surprised me, despite the hypothesis in our previous work on this topic.

While participants with combined technical and policy expertise were identified as more common in this process, multiple people I spoke with also noted that this may not have been evenly distributed; while industry organizations may have generally been more resourced for participation, technical expertise and familiarity was more likely to be present among privacy advocates and among the more traditional Web standards participants.

In addition to individuals who bridged technical, legal or business expertise, Do Not Track and other multistakeholder processes have often brought together people with disparate educational and professional backgrounds. My interviews are peppered with the informal comparisons that business executives make about technical people, or that engineers make about lawyers, etc. While these (over)generalizations may be interesting, I'm not certain how valuable they are to report here. More relevant to questions of participation and what determines success in multistakeholder process, though, is the challenge and importance of communication between people with very different career backgrounds. For example, see the success criterion of learning as part of [convening in the standard-setting process](#).

5.5.6 Conclusions for legitimacy and efficacy Participants in technical standard-setting processes for developing Internet and Web protocols are certainly not demographically representative, of the world, of the user population or even of high-income Western countries. To the extent that relatively easy access to participation could provide procedural legitimacy, multistakeholder processes may have some advantages, but these standard-setting bodies still fall far short of statistical representation. However, it's possible that certain policy-relevant areas, including accessibility and perhaps also privacy, may have more parity on one demographic dimension (gender) – this is worth further study, but may indicate either that especially values-oriented topics are more likely to attract a broader range of participants or that some sub-fields have more proactively welcomed broader participation, perhaps because the legitimacy of diverse participation is recognized as important.

Beyond demographics, we often see legitimacy of a process by how stakeholders participate in decisions that may affect them. How stakeholders are defined may influence how these multistakeholder processes function and the “two sides,” industry/advocacy, implementer/non-implementer perspective is commonly held, and either a symptom or a contributing cause of entrenchment. Recognizing the complex, multi-sided arrangements of Internet and Web services may help in

identifying promising techno-policy standards work.

Finally, participants are, we must remember, individuals, not just representatives of organizations, and the roles, backgrounds and relationships they have influence how multistakeholder processes operate. Leadership, not mere moderation, in formal or informal ways from prominent and invested participants can be a driving force and has been especially significant for security and privacy. Where technical standards have particular impacts or interactions with public policy, there is a value for individuals who have both technical and legal expertise, and an apparent trend towards participation by those multi-disciplinary professionals. While that may be a shift in the background of participants, nonetheless many involved in Do Not Track had little or no previous experience with technical standard-setting or the rest of that community. Cross-boundary communication and collaboration is a potential boon of techno-policy standards, but the lack of close connected ties present in existing technical standard-setting communities also demonstrates the challenge of building effective working relationships.

5.6 How participants see privacy

In my conversations with standard-setting participants, I asked about their own views on privacy: how they considered it as part of their work, in their lives personally and for users of the Internet. Given the important decisions these engineers, protocol designers, lawyers and executives make regarding online privacy, it's important to understand their own perspectives. It also served as a useful introduction into more specific questions about Do Not Track or experiences with particular technical standards that might have privacy impacts.

As described previously,¹⁶³ privacy and security are values of special relevance to the Internet and the Web. But though they are distinctly relevant, security and especially privacy are still complex and contested ideas and their application to the Internet or to software engineering in general is not settled.

Here I will show the range of views of privacy as a concept to assist in understanding how those mental models affect decisions about privacy on the Internet. Next, I look at some common touchstones that drive motivating examples for technical standard-setting participants, including particular sensitive datatypes and their implications. One touchstone in particular – how one thinks about privacy for one's own children – helps illuminate the ways that participants think about privacy for others. Finally, tied to these different concepts of privacy and thinking about privacy for others, participants speak about the actual work of privacy in technical and legal settings.

5.6.1 Views on privacy differ It's common to talk about privacy, a notoriously complex, challenging and challenged concept, through some narrower property, goal or sense. With the people I talked to, it often seemed that someone would start talking about privacy-as-something – and even for people who explicitly recognized variations in their own views on privacy and variation in the views others hold (see below), it would frequently be useful to talk about one particular sense or part of privacy at a time. Privacy-as-x can include a wide variety of concepts: the privacy analytic from Mulligan and Koopman identifies 14 distinct dimensions for classifying claims of privacy (Deirdre K. Mulligan, Koopman, and Doty 2016), which I'll refer to regularly.

5.6.1.1 Privacy-as-confidentiality An easily accessible example is privacy-as-confidentiality: a sense of privacy as keeping certain facts secret from others. This

¹⁶³See *Privacy and Security: Values for the Internet*.

is brought up occasionally by interviewees, often as a contrasting concept, to say that others used to see privacy in this narrow way, but they realize that it's more than that.

Here privacy is described as protection against the threat of violating confidentiality:

privacy would be something related [to security] but a little bit different. It'd be more of like [...] trying to discover something that the user thought was hidden but is really not.

But confidentiality can also be invoked as a historical contrast:

So 20 years ago when we thought about privacy it was really secrecy, right. It was, "Don't tell anybody anything about me"

Researchers have also noted this distinction, that privacy-as-confidentiality may have been an early attempt at privacy to engineering, as driven by applying cryptographic functionality from security engineering (Danezis and Gürses 2010).

5.6.1.2 Privacy-as-control More commonly accepted or invoked was some sense of privacy-as-control, what I would categorize as either informational self-determination (A. F. Westin 1967) or a sense of user understanding and capability to effect a choice about information. In the privacy analytic, control over personal information is the *object* of privacy, it's what privacy gives you.

From a very software engineering perspective, this gets described as permission, consent and control over software:

when you look at a lot of web APIs you have to make sure that you're keeping the user first, and I think that's the mantra that we tend to talk about, making sure that the person using our products is in a position where they can make knowledgeable decisions about what they want to do with software.

To me, the most useful [definition of privacy] is the right of an individual to control what happened to his own information. And it means that I may decide voluntarily to give certain information to another party. [...] I may decide voluntarily I'm willing to give it to this party in exchange for whatever benefits I derive from it. I tell Google my current location, it gives me a map of the area, the stores

around me or something. I know why they want to benefit from it but I'm getting some benefit too, so I will do it.

From a business perspective, control may be closely tied to transparency and the availability of controls. We might see these either as distinctive views of the concept of privacy (e.g. that privacy is the ability to see what information is collected and what controls are available) or as combining both what privacy is and what functionality has to be implemented for the value to be maintained.

Privacy for us was primarily the interaction with a consumer and how information was either collected or what controls were provided to them and what disclosures, transparency came along with that.

It's clear that the different models of privacy have overlaps and connections. Transparency frequently comes up in the context of privacy-as-control because how can someone meaningfully exercise control if they don't know what's happening?

We're very clear about the information we get through [data source], what we do with it, what we don't do with it. We say what we don't do with it. And so, people can make that choice.

5.6.1.3 Privacy-as-protecting-data Related to privacy-as-control but with less focus on the user interaction are senses of privacy regarding data handling or how data about a person is used after it's collected; these might be defined as the *target* of privacy, the personal data that is being protected (Deirdre K. Mulligan, Koopman, and Doty 2016). That could be as simple as not publishing log files or more complex enterprise privacy management systems.

So usually, for me, privacy means you have personal data: it could be IP addresses, it could be email, it could be whatever. And privacy research is about how to best handle this data, protect the data, make sure that the data is used according to consent.

I think there were certain types of activities that [Company] felt like it would like to be able to do if it took reasonable steps to protect data. And there were people at [Company], and many of them product or engineering people, who were very, very cautious. They would call up all the time, can I do this? I've put a flag on this. I have this data separate over here. There were a lot of people really taking care, but

within the context of taking care and pseudonymizing data there were also tremendous business pressures.

These senses of privacy in data handling often have some sense of responsibility, stewardship or appropriateness about how data is stored or used. From a European regulation perspective, these concepts might be more familiar as much privacy-related law is specifically about data protection (typically, ensuring conditions for processing of personal data), rather than privacy rights.

5.6.1.4 Privacy-as-context-sensitivity While less commonly raised than these concepts of control or protecting data, other conceptions of privacy were significantly identified. Related to personal information but touching more on social norms and appropriateness would be respect for context. Most interviewees are not specifically referring to Nissenbaum's theory of contextual integrity (2004), but may be influenced by it; there are multiple sources of "context" as a source for privacy in engineering, including ubiquitous computing (Benthall, Gürses, and Nissenbaum 2017).

How one tech employee described context and appropriateness:

Another aspect I think that we miss is that we have personae in real life, what I do at work and what I do in my hobbies and what I do at home or what I do in voluntary work and for other people, these are all distinct aspects of myself. If I volunteer, I'm going to pick an obnoxious example, at a clinic for sexually abused children, right, I do not need adverts about sexually abused children following me at work, and so this happy way that the online services just mashes all together, you know if my job at work is doing video quality assessment of online videos and some of them are pornography that doesn't mean I'm interested in pornography at home. So that's another aspect of online privacy that I think we completely missed, that is it appropriate now, is this contextually appropriate, and that's privacy again.

In this quote and a few other conversations, there's an identification that some information is specific to a particular situation or part of life and not appropriate to come up elsewhere, what we might call the collapsing of contexts as in danah boyd's work (2008).

Distinct but related are some ideas of privacy-as-relevance, that your privacy can be violated by "too much information" or information that you didn't want to come across about family or colleagues.

5.6.1.5 Privacy as freedom from intrusion While still related to information about people, privacy-as-relevance or privacy-as-context-appropriateness connect to the privacy concept as freedom from intrusion. To connect to the privacy literature, we would typically look further back, to Warren and Brandeis and ‘being let alone’ (1890), a particular interest in the late 19th century when photography and newspapers were technologies changing the basic assumptions about intrusions into our daily activities.

Two participants particularly highlight this idea of being left alone in the context of targeted advertising, expressing feelings of annoyance.

And as long as they’re left with the opinion that users don’t care they’ll do whatever flashy thing makes the most awesome user experience where, you know, you buy shoes online and they deliver special shoelaces to you in the next day because they think you’re awesome and they think that you want that. Some people do, <laughs> you know. I don’t. <laughs> I want you to just go away. I don’t want to have any interaction with these people. I just want the thing that I ordered, you know.

I used to think that personalized advertising would be an improvement over general advertising, but actually I find it hugely annoying and intrusive, and it’s stupid in many cases, you know, I wasn’t looking at that for myself, I was looking at that because my friend Nick was in my office and he said, “Maybe we could find a product online,” I was looking for my son, I was looking for him, or I’ve already bought the damn thing, have you not noticed I’ve already bought the damn thing, and so the way it follows you around, it’s sort of like having a terrier, it’s constantly going, yap, yap, yap, behind you all the time, nipping at your heels, it’s just infuriating.

Analogies to the Do Not Call program in the US have been familiar in Do Not Track discussions (where the Do Not Track name comes from), despite rather large differences in design and implementation. Also, Do Not Call is more narrowly targeted to privacy in the sense of intrusion (telemarketers ringing your landline during dinner), although intrusion (in addition to concepts of control over information) is also sometimes identified as relevant to online advertising.

Intrusion (for example, the physical intrusion of stalking) can also be a frustration to maintaining one's own autonomy, a value identified as protected by privacy.¹⁶⁴

And I kind of never felt that autonomy even as an adult because I was then growing up with the internet, and so as an example, I went to my boss and said [...] "Okay. You're putting our work schedules up where anybody can see them, and I have somebody showing up at my place of work before I get there," and it was "Well, too bad. We're not going to change what we do," and they were online, and that's just how it was.

5.6.1.6 Recognizing the variety of views of privacy The variety of senses of privacy that get discussed also reflect varying levels of concern among participants about their own privacy. And that variety among the population is something the participants themselves recognize about Internet users, which has important implications for the design of Web technologies.

For one baseline characteristic metric of privacy concern, I surveyed interviewees based on the Privacy Segmentation Index,¹⁶⁵ which divides people into the categories of privacy fundamentalists, privacy pragmatists and the privacy unconcerned. That index has been used to show general trends in the public: that many (and a growing number) are pragmatic about privacy, while people who are unconcerned about privacy shrinks as a proportion (perhaps because of increased awareness) and privacy fundamentalism is a growing minority (Kumaraguru and Cranor 2005). Among my interviewees, only a single one was classified as unconcerned, the majority (16) were pragmatists, and a substantial minority (8) were fundamentalists. (It often was not obvious to me, even among people I know professionally, what category an interviewee would fall into.) This generally reflects the trend in the public index, but our group of technical experts, privacy lawyers and advertising industry employees are especially unlikely to be unconcerned or unaware of privacy.

And while personal and professional perspectives on privacy certainly vary among my interviewees, participants also recognize or conclude explicitly that views of privacy differ among different professions, cultures, age groups and especially among the body of non-expert users of the Internet and the Web.

¹⁶⁴ Again following the privacy analytic, autonomy may be a *justification* for privacy, a reason that privacy is needed.

¹⁶⁵ Also described as the Core Privacy Orientation, see A. Westin (2001).

I think most people are somewhere in the middle and they have, you know, different things that they post that want to go to different audiences or they want everything to go to a somewhat larger group of friends. But I don't think there's a one-size-fits-all. I mean, I think about it much more in terms of letting people understand what's going on, letting them make choices that are right for them, rather than us deciding, you know, everything has to be public or everything has to be secret.

That views of privacy differ substantially among users is one core reason to pursue user choice or preference expression mechanisms at all. Without such a difference, added infrastructure to enable choices and communicate preferences would be unnecessary intrusion: if tracking of online behavior is harmful or always unwanted, then blocking it is more efficient and beneficial than letting users make a choice about it; if tracking of online behavior isn't a genuine privacy concern, then letting users choose not to be tracked wouldn't provide any advantages. This conclusion is a key motivation behind Do Not Track and other expressive privacy features: regarding the different paradigms possible,¹⁶⁶ DNT provides the end user with a variety of choices that are then communicated on to participating parties, rather than relying on a singular view of privacy interests.

A more complete list of the privacy-as- concepts identified in my corpus is included as [an appendix](#).

5.6.2 Touchstones for privacy and impacts on others How do participants in technical standard-setting talk about privacy in their own lives or in designing online services? Rather than falling back on abstract, philosophical language, it was very common for people I talked with to jump to particular motivating examples, whether it was specifics about their own life or hypotheticals. While the range of those touchstones was wide, particular topics were often repeated, especially sensitive datatypes (regarding location, sexual orientation or health) and family members, especially their own children.

5.6.2.1 Sensitive datatypes and salient privacy topics One direct way to conceive of privacy and explain its importance is to focus on the particular *target*

¹⁶⁶See, previously, [Do Not Track](#), a “handoff”.

of privacy, on what it is that we think privacy is meant to protect (Deirdre K. Mulligan, Koopman, and Doty 2016). While it was common for participants in technical standard-setting to refer to views of privacy as control over information, they also identified the particular datatypes over which control were important, either to them or to the users they think about.

Several participants in IETF and W3C technical standard-setting referred to the privacy implications of geolocation functionality – that a device or online service can determine (with sometimes uncanny precision) where you’re currently physically located. Location has particular salience for privacy because of three distinct properties of location data:¹⁶⁷

1. it reveals other information (health conditions, employment, social connections, etc.) about people, based on where they go;
2. it’s often uniquely identifying;
3. it facilitates physical intrusion.

One engineer discussing the Geolocation API directly touches on (at least) two of those factors:

We don’t want to give any information out that we don’t absolutely have to. Location is a very sensitive one where if you travel from your house to work every single day, the service provider is gonna have a pretty good idea of where you live. In fact, if a service provider sees you going to a 7- Eleven instead of a Peet’s Coffee they can make decisions about your lifestyle and what economic status you’re at

However, part of why geolocation privacy in particular is raised so frequently when talking about technical standards is that at both IETF and W3C, defining APIs for communicating precise geolocation information was one of the first experiences with mobile device sensors, and the debate and architectural models would become the basis for many subsequent technologies (camera, microphone, light sensors, accelerometers, fingerprint readers, and on and on). Interactive user permissions on the Web started with Geolocation, and there were (relatively) heated debates over sticky policies (user’s being able to specify machine-readable permission about use and retention) between IETF and W3C.

¹⁶⁷ Alas, as an impatient scholar, I’ve been presenting this three-part framework since 2010 without formally publishing it. See slides (Doty 2010) and related report (Doty, Mulligan, and Wilde 2010).

Other sensitive datatypes cited include health information, or particular categories of health that seem especially sensitive. As someone in the ad industry described it, advertising based on certain sensitive topics themselves seemed bad for societal outcomes:¹⁶⁸

it's a little problematic because there's no definition of "sensitive" [...] but what I was mostly concerned about, and it ties back to the other one about chilling effects – mental health, for example. Companies create very sensitive interest profiles on mental health in ways that I personally didn't think was a great thing for industry or society, and we decided that's sensitive, right?

While it's acknowledged that "sensitivity" of information is difficult to describe (perhaps in much the same ways that "privacy" is), a connection is made to chilling effects – that knowing that sensitive information about you is collected and used might discourage you from learning or discussing those topics that are sensitive to you. Sexual orientation was raised by multiple participants as a sensitive datatype regarding interpersonal relationships, but also in the context of a fear of inhibiting discussion or chilling young people from learning more about sexuality.

5.6.2.2 Privacy impacts for others Privacy is a sensitive, personal, subjective, contested value, which motivated my asking standard-setting participants – people who debate and design protocols that implicate online privacy for Internet users – for their personal views on privacy. But the participants in these interviews, and the participants in technical standard-bodies worldwide, and the employees of tech companies that build online services, are in many ways not similar to or representative of the population of users of the Internet. Based on demographic categories but also based on technical savvy or knowledge, the developers of Internet protocols and software are quite distinct from the median end user.¹⁶⁹ It

¹⁶⁸The particular limitation here is on the *use* of these sensitive categories rather than their *collection*, so it might be that the target of privacy is not specifically the data itself, but harms related to targeted messages about people within those sensitive categories. However, the sensitivity of use may also be related to potentially disclosing a sensitive health condition to others based on the presence of targeted advertising on that person's device, in which case we might say that the *target* is the personal data about health conditions and the *from-whom* is friends, family or people who might share a device with the individual. Having clear orthogonal dimensions can make it easier to tease out these differences.

¹⁶⁹See [Who participates and why it matters](#).

is perhaps as important then to consider what designers think about privacy for other people as they think of privacy for themselves – the *subject* of privacy in the analytic mapping (Deirdre K. Mulligan, Koopman, and Doty 2016). While I included a prompt in my interview guide to uncover ideas about user thoughts on privacy, it often came up unprompted, in three ways:

1. distinguishing that the speaker was not concerned about their own privacy, or that the speaker recognized they were more concerned about their own privacy than others might be;
2. noting the lack of understanding by users of the Web about how technologies that affect online privacy work; and,
3. identifying family members as a particular and compelling case of concern for the privacy of other people.

Why might these interviewees not be concerned about their own privacy despite their knowledge and work in a privacy-relevant field? For one, because the participants in the technical and legal fields tend to have many advantages and privileges of class, race, educational background and (relatively) stable governance.¹⁷⁰

That's just my personal interest. Because certainly those photos [of drinking in college] would have existed, and probably do exist in a Polaroid somewhere. But there's not a lot of downside there. I personally am not terribly worried about government data collection about me. I understand why people are. And I tend to be more trusting of certainly the U.S. government from – not because I think that they're adept at protecting privacy or data, I just don't think that they're nefarious, and I don't have much – I don't really have anything to hide. And so that doesn't really worry me. So I think if they can be subject to similar baseline requirements like data security, then it doesn't worry me that much.

This form of explanation – the lack of risk (“downside”) and the lack of “anything to hide”¹⁷¹ – emphasize how the lack of concern about personal privacy in

¹⁷⁰These advantages and stabilities are described further in directly considering the ethical implications of “studying up” around this population.

¹⁷¹Writing on “nothing to hide” as a fallacious argument is widespread and I wouldn't be sure who to cite on the rhetorical topic. I don't take this individual's passing remark as an endorsement of a “nothing to hide” argument against privacy as a value of importance and I don't include it as

these certain threat models is contingent, and the speaker repeatedly interleaves the explicit idea that these are personal calculations and will be different for others.

Some interviewees are also less concerned about keeping things private specifically because their own work is done in public or might involve some publicity. That can range from people who consider themselves public figures to engineers who just do more work online:

But for example, so I'm in a gym and we have lots of events and so on there and when they send out emails, oftentimes they'll send an email and they'll have like a long CC list and I always react, "That's not really cool because some of these people might not have wanted their email address shared." I personally don't care, I mean, my email address is super easy to find and this is a pretty common way to react, I don't personally care about a lot of these things but I am very aware I think about when people's private information is shared.

Despite the relative privilege and advantages that people I spoke with share, some also identify themselves as in some cases likely to be more concerned about their own privacy than others.

So part of why I don't use Facebook and Uber and LinkedIn is because of their track record with what they do with information, and there's a real cost to your life, right? I was in [City] on Monday, and it took me about a day to realize that's a town that no longer really has a functioning taxicab system. Apparently it was a weak system to begin with, and it got just decimated by Uber and Lyft, and it was so bad that I downloaded Lyft Monday night and used it to get around town on Tuesday. They were my first and second Lyft rides ever, and this is after three or four or five years of everyone in the world telling me that, "You can't function in human society without these apps." So that's one example [of things that might seem unnecessarily paranoid to others]. I mean, that's justified paranoia, but that's one example.

[...]

a criticism of that perspective. Indeed, one of the primary reasons that nothing-to-hide is a poor argument against supporting privacy – that privacy is a value for protecting people who may have less power or protecting society so that people can take unpopular positions – is demonstrated by an individual distinguishing their personal fears from others'.

Well, two billion Facebook users can't be wrong, right? So I'm not trying to make any super-nuanced points about empirical research I've seen. I think I'm just reflecting on an increasing feeling that my choices are out of lockstep with just about everyone I know personally and also with what I read in the press about what the world is doing.

“paranoia” can be a term suggested to describe being outside of a community's social norms, rather than its formal denotation about irrationality. And in contrast to the privilege distinctions, these different evaluations of privacy can be among people who are similarly situated (“everyone I know personally”).

One theme that gets at that kind of distinction – where others might not be concerned about themselves as subjects of privacy while others identify it as a concern – is user understanding, or more often the lack thereof, about Web technologies and their privacy implications. Some of these assessments are quite blunt:

So I had done usability research, and I understood that people were by and large clueless about where their data was going

Users don't know what companies collect this information about them:

I was fired up about it. I still am. The notion that a company I've never heard of has a list of websites I've gone to is not awesome, and I think folks – actually, I think there's plenty of science showing folks don't like it and would like to be able to limit it, and so I was concerned about it.

Users only understand when triggered by a particular event¹⁷² and user attention and understanding are hard to persist over time:

most of the time, of course, unless something happens like that, triggering, what the fuck, you know, how do you know that I know these 40 people, unless there's a triggering event like that, of course, most users don't notice, and it takes a disaster for them to notice, and, of course, we don't want to run the industry such that we run until we

¹⁷²There's a separate code in my dataset on “exogenous events” which I initially anticipated to be about Snowden revelations, which do come up in that sense, but Cambridge Analytica is also frequently cited.

hit the iceberg and then we panic, we'd rather not hit the iceberg in the first place, thank you very much, but I have a fear that we're going to hit the iceberg.

That's the problem with this stuff. Unless you are constantly reminded of it, you forget about it, right? That's the general mass of people on the web. They get pissed that Facebook put some complex thing to read that they know is not improving their privacy but taking it away. They get pissed for a day – whoosh – and they're right back in their normal life. They don't change. They don't jump out. The problem is the threat is not – what's the word – acute, right? It's gradual, and so it's going to get you later in life kind of like before. It's not something you react to in the present tense.

And that users' lack of understanding or awareness or ability to control may be an intentional design outcome:

you can articulate that you care, but you have so much going on that it's really hard for you to take steps, which is why I would hope that the government would address the more significant harms, because people can't possibly understand, and that's intentional. I mean, that is absolutely intentional. Industry doesn't want them to understand. It's confusing. [...] you understand, asymmetric information: you can't grasp it. Even as a parent now, I deal with parents all the time. [...] Parents have no concept of what's going on.

In these quotes, interviewees connect the lack of understanding – because it's “confusing” or “gradual” – to the lack of taking action to prevent privacy harms (in these cases, typically collection of information about them). There is an implicit response here to the well-known “privacy paradox” – if users are concerned about their own privacy, why don't they alter their behavior more often to better protect it? Experts identify a lack of understanding in users, which provides an explanation of the lack of action.

5.6.2.3 Privacy for one's children Most surprising for me¹⁷³ in these interviews was how frequently people I spoke with cited their own children in describing how

¹⁷³Nota bene: I am not a parent.

they thought about privacy in their own lives. In part this may be paternalism in its original sense, that parents make decisions for their children because children may not have the awareness, understanding or knowledge to decide about information about themselves. But interviewees also recognize the lack of autonomy that children may have when parents are making choices on sharing information about them. It seemed that there was often more salience to the protection of children, the risks for their future lives and their ability to decide, than for the (privileged) parent themselves.

personally I think I am always aware of privacy-related issues when using the Web, right, in different contexts. So, for instance, if I were going to share ... I think everybody has rules about how they share data and how they share things on social networks, for instance. I don't generally share pictures of my kids or use their names when I'm writing stuff on Facebook, for instance. That is a personal sort of set of rules that I've hit upon. I know other people don't abide by those, and it kind of is a good example I think of how people have different views about privacy when they're using social applications in particular. To me the privacy issue isn't so much my privacy. It's about that if I'm sharing information about my kids they're not old enough yet to be able to make that decision in an informed way, and I don't feel like I can make that decision for them, so therefore I'm not sharing information about them.

Children and family may also be cited as a contrast, where you might not care about limiting your own public image but of course wouldn't make the same decisions for children, again with the connection to making one's own decisions:

By no means am I a private person. [...] However, of course, there are things that I don't wish to share with the world, or maybe I wish to share them certain audiences, but not others. My family is not as eager to be, you know, super-visible, so, I keep them from- I don't share a lot of pictures about my kids. My wife never wants to be shared or tagged. So, I look at a goal- and I don't look at privacy- and I argue that for most people privacy is not an absolute goal. We want autonomy. We want freedom to make decisions.

Or a contrast in terms of generations and how younger people might not appreciate the risks of sharing information:

my personal view of privacy, it's gonna make me sound like an old man. I worry that people younger than me don't realize how dangerous putting something up on Instagram is or putting something up on Facebook is, and I think they're probably – in society there are probably lots of examples of, “oops, I shouldn't have shared that” and some of the ramifications.

Considering one's child's privacy can also have an impact on how one thinks about their own privacy, out of the same basic concept of protectiveness and importance:

I just think, you know, it's probably social pressure. My wife puts pictures up of our kids, and so my kids are online, so why would I not put myself – I mean I'm certainly not as important as my children, right, so I think that may have had a large part to it.

There is a universal quality here about a parent's responsibility for, protection of and respecting the future choices of children.

5.6.3 The work of privacy A distinct way to talk about privacy is to talk about the work that “doing privacy” consists in.

5.6.3.1 Privacy-as-compliance Many interviewees (especially lawyers and less often people in engineering) discussed privacy in their work as privacy-as-compliance: less about the value itself and more about ensuring compliance with a privacy law or with a set policy. Many privacy teams in tech companies report up to the general counsel rather than through the product part of the organization. Or the privacy team is the “keeper of the policy structure” including laws and other negotiated agreements. This can have a few distinct senses though (and interviewees will often refer to more than one): where the goal of privacy work is to comply with privacy regulation; where privacy work is about maintaining internal or external accountability that policies and practices are being upheld; or, where privacy work is managing risks, which could be security breaches, or more downstream, the unwanted news coverage or regulation that privacy issues could spur.

So privacy is a great example where sovereign entities have laws and regulations in place on the topic, but those laws and regulations tend not to be written in such a way that they're immediately obvious

how one would implement those things. And frequently in order to verify whether or not people are meeting those regulations, there's a desire to see certification in some form of compliance that might be ascribed to those behaviors. And so in order to do that you have to have some kinds of controls that you put in place, as well as criteria by which those controls would be executed. And so we focus on coming to international agreements on those topics relative to large-scale regulatory requirements, or to establish foundational concepts in emerging areas, where we know that that type of activity is likely to happen.

All of those [businesses] have completely different perspectives on this concept of privacy. Some people think of it as compliance to a strict regulation, EU Safe Harbour or COPPA. Some people think of it as compliance to best practices [...].

5.6.3.2 “You have to sort of make it up as you go along” In addition to privacy work as a legal effort to maintain compliance with some external law or requirement, there is a distinct effort in the legal work of making internal, organizational policy to apply to some technical or business practice.

there are areas of grey, right, where we don't have an established policy, we're looking at doing something new or novel, and therefore we can provide guidance to the organization to say this is our policy area, we don't have a policy, let's say, in your specific area, but here's where we would say the risk profile is for this particular area, and then we would give our recommendation on where they want to go. In those scenarios it's more of an assistive role in the organization.

And there is frustration with the de facto perspective of privacy-as-compliance in the professional sphere:

[Privacy is] about the ethical and responsible use of data about people. I don't view my job as compliance, which is the problem with privacy today.

Some identify privacy as less focused on legal compliance and more on policy development, in contrast to other legal work:

I just thought it was interesting and in flux, and it was clear that there weren't rules of the road yet in the US, so that's basically what I thought is that this is really interesting. And I was in a meeting where [other privacy lawyer] said a year or two ago, "So with GDPR, are we just going to become like other lawyers where we just follow the law?" And I was like, "Oh, my god. How boring would that be?" Privacy is not like that. <laughs> You have to sort of make it up as you go along – at least that's been the case in the past.

The "in flux" nature is attributed in part to the relevant youth of privacy in law and regulation, at least in the Internet context. But as a result, it makes the work of doing privacy as making it up, which might include lobbying, or interpreting new law, or arguing for policy approaches, as opposed to systems that just ensure compliance with more well-established regulation. That requirement for ongoing interpretation under broad or ambiguous regulation has been credited with empowering the field of privacy and bringing outside groups in to debate the privacy impacts of corporate actions (Bamberger and Mulligan 2015).

On the more technical side, there may also be a sense that the work of privacy is about discovery rather than simple implementation. Richmond Wong (2019) studies the field of human-computer interaction and explores how design practices can be used to explore, critique and present alternatives to privacy problems, in contrast to the perspective of privacy being a single fixed concept (like control) with design and engineering as putting that concept into practice.

Whether the work of privacy should be about contesting a particular concept of privacy is an open question. The argument that privacy is essentially-contested recommends that the "progressive competition" over the value is a beneficial feature that makes privacy more useful as a concept (Deirdre K. Mulligan, Koopman, and Doty 2016). But it's notable that in some cases the value or purpose of privacy might be obscured in how it's discussed or considered.

The tension between whether privacy is settled elsewhere (like through formal regulations) and then implemented vs. being contested in the same place that it's being realized recalls the tension between separation and integration in how ethical concerns more generally should be a part of engineering practice.¹⁷⁴ It also connects to competing notions of organization-centered vs individual-centered views of multistakeholder process.¹⁷⁵

¹⁷⁴See [Separation vs. integration](#)" in the earlier chapter on The Ethics of Engineering.

¹⁷⁵See the section of this chapter on [Individuals vs. organizations](#).

5.6.4 What to conclude from these diverse views of privacy Various privacy-as-control views are well-understood and common among this population of privacy experts and engineers developing technical standards that contribute to flows of information. That's no great surprise, but it should inform our understanding of the controls and mitigations that are likely to be considered in that setting. Different views of privacy, different threat models and concerns, may not get the same protection from additional transparency or data handling controls. How well will these views of privacy and corresponding expertise and developed tools and practices accommodate distinct privacy concerns: around fairness or online harassment, say? Or, as others have pointed out (Kostova, Gürses, and Troncoso 2020), how will views of privacy as control and control mechanisms work as software architectures change?

Recognizing different views of privacy means more than anticipating gaps during the engineering process. For compliance with privacy law and regulation, legal counsel are considering how to comply for Internet services that cross jurisdictional lines; for attracting customers from different countries and cultures, product designers are considering how to appeal to people with different cultural attitudes towards privacy. As privacy continues to be contested, there is an impulse to accommodate that ongoing debate with architectural designs that support public policy values without first settling all questions about their exact scope.

Understanding, effective capability and power are explicitly identified factors that respond to the motivating question about responsibility within the socio-technical system. Recall the vignette of “An ad that follows you”¹⁷⁶ where it isn't clear who is responsible or what you the user could do differently.

While tempting, we don't need to conclude that because experts in Internet protocols, online advertising and privacy draw a connection about the privacy interests of their children that privacy experts are advocating for a policy position of online paternalism. Nor should we conclude that paternalism is the proper or most effective approach we should pursue in looking at how to design for privacy among a non-representative group of end users.

Some conclusions we can draw from the significance of parenting as a theme, though. First, experts and designers of Internet protocols and online services may be attuned to thinking about the privacy interests of people different from themselves: many recognize the variation in preferences, levels of understanding and values about different conceptions of privacy. Second, there are mental models readily at hand for considering the impacts to people who are less expert or less

¹⁷⁶See *Do Not Track*, a “handoff” in the earlier chapter on Privacy and Security.

capable of making their own decisions – people are familiar with the privacy of other people and people who can't decide for themselves from their intimate lived experience in raising children. In addition to exploring inclusive processes, participatory design approaches and user research grounding, we can also identify that thinking about the impact on differently-situated others is an existing practice in the technical field of Internet privacy.

Finally, competing views of privacy are complemented by competing views of privacy work. When privacy is enacted in developing technical standards, is that the work of debating the concept of privacy and the normative questions of what we should protect or how responsibility should be distributed? Or is the work a more technical matter of reifying policy that has been decided elsewhere into concrete form?

5.7 Towards integration

At the [opening of this chapter](#), I outlined the two high-level research questions of this project and the five clusters of themes from my empirical findings that speak to those questions. Those themes have touched in different ways on the two research questions. But they also recommend a challenge and an opportunity for the larger justification of my project: how to better support values such as privacy through the techno-policy standard-setting process. Below I summarize the findings in relation to my research questions and the opportunity and challenge they present. In both cases, I see a common key, the deeper and more nuanced integration: of values into engineering work, of different kinds of expertise, of technocratic and democratic process.

5.7.1 An opportunity and a challenge Regarding the impacts of multistakeholder techno-policy standards-setting processes on resolving public policy disputes for the Internet:

Consensus-based multistakeholder technical standard-setting process provides a real opportunity for stable, cross-boundary collaborative solutions to disputes over public policy values in socio-technical systems. Those solutions, though, would require overcoming difficulties at several stages in the standard-setting process, under conditions where implementation and interoperability are well-aligned with substantive protection of values and accommodation of ongoing contestation. And under those conditions, we should still anticipate tension between representational vs collaborative views of the individual or democratic vs technocratic views of the process and heated conflicts from diverse or antagonistic participants.

To take advantage of this opportunity, I argue, we must embrace the integration of those representational and collaborative views and design processes to accommodate heterogeneous perspectives. As described in [Chapter 1](#), collaborative governance requires a problem-solving orientation and ongoing participation from stakeholders; findings show both the promise and the deep challenges to productive engagement among potentially antagonistic parties (5.2). As described in [Chapter 2](#), engineering is inherently ethically-laden and the engineering ethos is individual, practical and engaged; individual participation has proven to enable autonomy and principled contribution at the price of conflict over who one represents (5.4). As described in [Chapter 3](#), privacy is and will be contested and that contestation can be productive; we have seen that participants evaluate communication and learning, especially across disciplines, as an important success (5.2). Put together, opportunities to address values such as privacy in socio-technical

systems need multistakeholder processes where engineers are actively engaged in problem-solving, in learning and in negotiating with stakeholders.

Regarding standards-setting participants' views of privacy and the resulting impacts on Internet user privacy:

Findings on participation (5.5) emphasize the relevance of this inquiry into views of privacy from those who are designing Internet protocols and negotiating Internet standards. Standard-setting participants are not generally representative – in demographics, in level of expert knowledge, or otherwise – of the population of Internet users. Lack of representation presents a substantial challenge to the legitimacy and responsiveness of techno-policy standards in addressing privacy.

But a substantial challenge is not a lost cause. Privacy, accessibility and other areas of public policy interest may already attract relatively more diverse participation. Furthermore, we should proactively seek better ways to support privacy from our current systems of design and governance based on what we have learned about current participants. Standard-setting participants have widely varying views on the conception of privacy and directly acknowledge that views and priorities differ. Even more intimately, the parent-child relationship or other views of family present a touchpoint for considering privacy (and other values) for differently-situated others, and not just in the sense of paternalism but also in valuing the autonomy of others. The work of privacy is seen as simultaneously continuing to figure out privacy as well as realizing or stewarding it.

As described in [Chapter 1](#), both procedural and substantive legitimacy are important for governance and focus on interoperability and rough consensus will not be enough to assuage all concerns, particularly given non-implementer stakeholders (5.5). As described in [Chapter 2](#), the numerous detailed decisions of engineers can have a large impact on the deployed technology of the Internet and the Web and as we've seen (5.4), engineers have remarkable independence even from their employer in the positions they take in technical standard-setting. As described in [Chapter 3](#), privacy is essentially-contested and so debates over privacy will not be settled and should be considered both in concrete user needs and in high-level goals and technical architecture; in the findings, we have seen (5.6) that privacy views vary and are recognized as diverse and that interdisciplinary expertise – involving policy, ethics and technology – is especially valuable. To respond to the challenge of representation and the need for legitimacy in governance, addressing privacy, a value inherent to the social use of the Internet, will require increasingly interdisciplinary work – involving policy, ethics and technology.

5.7.2 Integration is key Key to the answers to both research questions is *integration*: of values into engineering, of different kinds of expertise, of technocratic and democratic process.

That integration is key, or that integration is worth pursuing as an opportunity, does not imply that the result is simple. These findings do not support a simplistic integration of the form of embedding or hard-coding an unchangeable value in a permanent, unquestionable or unaccountable piece of architecture. Similarly, they don't guarantee that multistakeholderism is a panacea that will guarantee integration of every diverse interest or perspective.

Rather, this suggests integrating the debate over values along with expert evaluation of technical design and integrating training and collaboration to encourage more professionals with tech, policy and ethical expertise. That integrated work and training can prepare us for the more holistic project of technology and the good life. But this nuanced integration should also accommodate diverse, conflicting participants and the impulse for separation and flexibility. Indeed, this is a hallmark of Internet standard-setting and the Internet's architecture: a contentious but collaborative development that supports common goals while maintaining diverse and flexible uses.

Handoffs are a theoretical tool for this more nuanced view of integration: handoffs are shifts in distribution of multi-actor responsibility in the context of a larger socio-technical system. Looking at Do Not Track as a potential handoff,¹⁷⁷ a new distribution is possible where a value of privacy is both integrated into a technical design but also enforced through a distinctive distribution of technology, regulations and norms. Values are often going to be integrated or embodied in technical designs one way or another, but we can choose how to intentionally enact the values we care about and design the form of their distributions.

In looking to [future directions](#), I will, finally, consider some possible direct interventions related to this promise of nuanced integration and suggest how to recognize future handoffs.

¹⁷⁷See [Do Not Track](#), a "handoff".

6 Directions for future work

What this leaves for the future is the question, or rather, the challenge, of what practices we could use in technical standard-setting to more effectively enact privacy and security for the Internet and the Web.

6.1 A triad for interventions

Throughout this research project and throughout my personal and professional efforts to support privacy on the Web, I have seen how potential improvements can involve three distinct but connected areas:

1. the people involved in the development of the technology;
2. the processes used in organizing its creation; and,
3. the tools used for design and implementation.

These might conceivably apply to questions of values in the design of technology generally, but I have observed them explicitly in the collaborative, rough consensus standard-setting process in particular.

6.1.1 People Leaders have played a substantial role in the support of security and privacy in Internet and Web protocols, perhaps especially because the standard-setting process doesn't rely on a single firm's hierarchical model but instead pushes for interoperability and collaboration between disparate and competing organizations. Leaders have provided both a backstop and a motivation for security and privacy to be considered more directly in the design of the Internet and the Web.

In recognizing the inherently ethically-laden nature of engineering, a shift towards integration of values in design and engineering and the potential for techno-policy standards that explicitly involve values such as privacy, particular combinations of expertise are increasingly useful. Those with both technical and legal backgrounds may be able to recognize and evaluate possible socio-technical configurations. While it may not be a rapid intervention, education can help meet this need. Schools of Information pursue an interdisciplinary approach, typically combining computer science topics with social science, law and policy and user-centered design.¹⁷⁸ Technology & Delegation, a seminar class, a lab class and a set

¹⁷⁸What an iSchool is remains an open question welcoming constant refinement and definition, but see for example the iSchools organization: <https://ischools.org/About>.

of curricular resources,¹⁷⁹ has been an explicit project to encourage students with varying backgrounds to confront direct intersections of technology and policy and how they interact.

The need for technologists engaged in public interest work and helping civil society and philanthropy has been described as a “pivotal moment” (Freedman et al. 2016). Scholars, foundations and practitioners have sought to develop a new field of public interest technology (Eaves et al. 2020), not unlike our earlier definition of a “citizen technologist” (Doty and Panger 2015). A network of universities is committed to “growing a new generation of civic-minded technologists” – an urgent and important goal.¹⁸⁰ Clinics provide students with experiential learning while fellowships directly integrate technologists into traditional policymaking spaces.¹⁸¹

6.1.2 Process While motivated individuals or community leaders have made a significant difference, organizational processes can bring broader and more systematic considerations of privacy, security and other values to the Internet standard-setting process. Rotating assignments in a Security Directorate at IETF is credited with improving the consistency of security reviews in Internet protocols, and similar attempts have been made with triggering wide review, including privacy reviews and architectural design reviews, at W3C.¹⁸² Procedural requirements can also be a hook for interested individuals to provide feedback on features that affect important values like privacy.

Clear and systematic process also provides an opportunity for more confidence in how consensus technical standard-setting can apply to policy-related topics. Removing uncertainty could remove confusion or even encourage cooperation. Along the same lines, we might ask for clearer roles from policymakers in their participation in consensus techno-policy standardization – how invested they are and what they aim to contribute, whether that’s requirements, some democratic legitimacy, incentives to participate or the power to enforce standards.

Finally, process implies or perhaps even requires continual application. Systematization, clarity, establishing roles – these would all benefit from repeated, ongoing processes that proceed to address the next tech policy and continually

¹⁷⁹Most recently taught in Fall 2019, with a wiki of Techdel resources.

¹⁸⁰<https://www.newamerica.org/pit/university-network/about/>

¹⁸¹See, for example, TechCongress: <https://www.techcongress.io/>

¹⁸²This theme was highlighted in Doty (2015a) and I believe systematization has slowly increased since.

review and revise existing systems. Periodic events, or development lifecycles that follow a linear waterfall model, don't provide the same opportunities for building relationships and effective institutions.

6.1.3 Tools Technologists must not forget the tools that influence tool-building. Tools here can range from simple, automated checks to comprehensive high-level design principles. Regarding security and privacy considerations in Internet standards, automated prompts can ensure that specification authors are at least aware of the need to directly address those values in new protocols. But simple, blunt requirements alone will also prove to be insufficient (Doty 2015a). Detailed guidance might prove fruitful, perhaps especially for those interdisciplinary or values-minded individuals who want to directly address privacy or security details in their domain of interest. I've tried to contribute for my part guidance on mitigating browser fingerprinting (Doty 2019), because it is a detailed privacy topic that accumulates across different features and could benefit from some coordinated and comprehensive response.

Tools may be most effective, though, when they work in concert with people and processes. Questionnaires, for example, allow experts close to a particular domain area but not necessarily trained on privacy or policy issues in general to help in identifying potential areas that may need further review¹⁸³ and collect details that a privacy expert who isn't intimately familiar with the domain can use in evaluating implications. W3C now sees widespread use of a self-review questionnaire for both security and privacy ("Self-Review Questionnaire: Security and Privacy" 2020) and a similar questionnaire is included in IETF's privacy considerations guidance (Hansen et al. 2013).

In the longer term, though, support for privacy, security and other values could be more efficiently maintained if they were designed in from the beginning, rather than spotted as potential problems along the way. Higher level design principles could be tools for these more fundamental changes, but privacy-by-design can be difficult to put into practice, even for those engineers who may already share the ethical commitment to it. Design patterns are documentation tools to codify and communicate abstract solutions to common engineering problems. Privacy design patterns, then, may:¹⁸⁴

- standardize language for privacy-preserving technologies,

¹⁸³This is sometimes called "issue spotting," inspired by the term from legal practice.

¹⁸⁴These project goals are taken directly from the collaborative privacypatterns.org project: <https://privacypatterns.org/about/>.

- document common solutions to privacy problems, and,
- help designers identify and address privacy concerns.

As a tool for communication, privacy design patterns can also facilitate communication of detailed engineering practice to lawyers or policymakers. Anti-patterns can help to classify the misapplication of a technique or warn of its unintended consequences (Doty and Gupta 2013) or to document the common problems that lead to a lack of privacy in Web standards (Snyder 2019).

6.2 Recognizing future handoffs

I have argued that privacy and security are values of distinctive salience to the Internet and the Web. But those concepts are complex, contested and likely to involve new senses over time. Even in the course of writing this dissertation, the distinctive, topical senses of privacy have changed. Fairness was a new privacy-relevant topic, with the idea that privacy might be the protection against unfair, society-wide inferences about oneself or one's community. Or perhaps privacy is freedom from the harassment and abuse that trolling and dog-piling on social media have made so easy. More recently still, the trend toward toxic disinformation uses those same social network channels to target not just an individual, but an entire society's sense of what is real or reliable.

Seen through the handoffs model, there are likely to be many more shifts in how values are maintained (or not) in different socio-technical configurations and how responsibility is distributed. Some paradigmatic shifts around security may be linear trends away from discretion or false reliance on assumptions of goodwill or end user expertise – like the ongoing march toward encrypting the Web. But there will also be the possibility of handoffs to more distributed approaches that involve communication among people, technology and regulatory systems.

Technical standard-setting – or specifically what I have called techno-policy standard-setting – provides an opportunity for multistakeholderism's promise of democratic and technocratic advantages, in the line of new governance as well as the bridging property of boundary organizations. Standard-setting's practical focus on interoperability suits it for handoffs to cooperative configurations developed by diverse parties – if those various organizations have incentives to pursue it and that heterogeneous group of individuals can work together. The handoff model encourages holism and asks us to look at the broader socio-technical system and the network of actors involved. Any multistakeholder process takes place

embedded in the context of ongoing technical, social, organizational and policy changes that influence it.

Whether these concerns are all considered senses of privacy or not, we face tech policy issues that are urgent, complex and have large impacts on public policy, including criminal justice, equal access to digital public fora, democracy and public health. We need comprehensive responses that integrate technical expertise, policy details and ethical understanding. To respond effectively and promptly, we must use what we have learned from our attempts to enact privacy on the Internet.

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>.
- Alvestrand, Harald T. 2004. “A Mission Statement for the IETF” RFC 3935. Request for Comments. RFC Editor. <https://rfc-editor.org/rfc/rfc3935.txt>.
- 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=>.
- Arendt, Hannah. 1958. *The Human Condition*. University of Chicago Press.
- Arkko, Jari, Brian Trammell, Mark Nottingham, Christian Huitema, Martin Thomson, Jeff Tantsura, and Niels ten Oever. 2019. “Considerations on Internet Consolidation and the Internet Architecture.” Internet-draft draft-arkko-iab-internet-consolidation-02. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/html/draft-arkko-iab-internet-consolidation-02>.
- Bahajji, Zineb Ait, and Gary Illyes. 2014. “HTTPS as a Ranking Signal.” *Google Search Central Blog* (blog). August 7, 2014. <https://developers.google.com/search/blog/2014/08/https-as-ranking-signal>.
- Bamberger, Kenneth A. 2006. “Regulation as Delegation: Private Firms, Decisionmaking, and Accountability in the Administrative State.” *Duke LJ* 56: 377.
- Bamberger, Kenneth A., and Deirdre K. Mulligan. 2010. “Privacy on the Books and on the Ground.” *Stan. L. Rev.* 63: 247.
- . 2015. *Privacy on the Ground: Driving Corporate Behavior in the United States and Europe*. MIT Press.
- 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.
- Benthall, Sebastian. 2015. “Testing Generative Models of Online Collaboration with BigBang.” In *Proceedings of the 14th Python in Science Conference*. <https://>

- [//conference.scipy.org/proceedings/scipy2015/pdfs/sebastian%7B_%7Dbenthall.pdf](http://conference.scipy.org/proceedings/scipy2015/pdfs/sebastian%7B_%7Dbenthall.pdf).
- Benthall, Sebastian, Seda Gürses, and Helen Nissenbaum. 2017. "Contextual Integrity Through the Lens of Computer Science." *Foundations and Trends in Privacy and Security* 2 (1).
- Berners-Lee, Tim. 1992. "The World Wide Web Project." November 3, 1992. <http://info.cern.ch/hypertext/WWW/TheProject.html>.
- . 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.
- boyd, danah. 2008. "Taken Out of Context: American Teen Sociality in Networked Publics." SSRN Scholarly Paper ID 1344756. Rochester, NY: Social Science Research Network. <https://doi.org/10.2139/ssrn.1344756>.
- 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>.
- Braman, Sandra. 2012. "Privacy by design: Networked computing, 1969–1979." *New Media & Society* 14 (5): 798–814. <https://doi.org/10.1177/1461444811426741>.
- Bray, Tim. 2012. "On the Deadness of OAuth 2." *Ongoing* (blog). July 28, 2012. <https://www.tbray.org/ongoing/When/201x/2012/07/28/oauth2-dead>.
- 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-w3c-is-a-restaurant/>.
- Burrell, Jenna. 2009. "The Field Site as a Network: A Strategy for Locating Ethnographic Research." *Field Methods* 21 (2): 181–99. <https://doi.org/10.1177/1525822X08329699>.

- Cargill, Carl F. 1989. *Information Technology Standardization: Theory, Process, and Organizations*. Newton, MA, USA: Digital Press.
- Caro, Robert A. 1975. *The power broker: Robert Moses and the fall of New York*. New York: Vintage Books. <http://www.amazon.com/The-Power-Broker-Robert-Moses/dp/0394720245>.
- Carruthers, Bruce G, and Terence C Halliday. 2006. "Negotiating Globalization: Global Scripts and Intermediation in the Construction of Asian Insolvency Regimes." *Law & Social Inquiry* 31 (3): 521–84. <https://doi.org/10.1111/j.1747-4469.2006.00022.x>.
- Cerf, Vinton, and Robert Kahn. 1974. "A Protocol for Packet Network Intercommunication." *IEEE Transactions on Communications* 22 (5): 637–48.
- Chung, C., A. Kasyanov, J. Livingood, N. Mody, and B. Van. 2011. "Comcast's Web Notification System Design." RFC6108. RFC Editor. <https://doi.org/10.17487/rfc6108>.
- Clark, D D, and D R Wilson. 1987. "A Comparison of Commercial and Military Computer Security Policies." *IEEE Symposium on Security and Privacy* 0: 184–94. <https://doi.org/10.1109/SP.1987.10001>.
- Clark, David D., John Wroclawski, Karen R. Sollins, and Robert Braden. 2002. "Tussle in Cyberspace: Defining Tomorrow's Internet." In *Proceedings of the 2002 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, 347–56. SIGCOMM '02. New York, NY, USA: ACM. <https://doi.org/10.1145/633025.633059>.
- Cohen, Julie E. 2012. "What Privacy Is For." *Harv. L. Rev.* 126: 1904.
- Coleman, E. Gabriella. 2012. *Coding Freedom: The Ethics and Aesthetics of Hacking*. Princeton University Press. <http://www.amazon.com/dp/0691144613>.
- "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>.
- Danezis, George, and Seda Gürses. 2010. "A Critical Review of 10 Years of Privacy Technology." *Proceedings of Surveillance Cultures: A Global Surveillance Society*. http://www.researchgate.net/publication/228538295_A_critical_review_of_10_years_of_Privacy_Technology/.
- Davies, Charlotte Aull. 2012. *Reflexive Ethnography: A Guide to Researching Selves and Others*. Routledge.
- Davis, Michael. 1991. "Thinking Like an Engineer: The Place of a Code of Ethics in the Practice of a Profession." *Philosophy & Public Affairs*, 150–67.
- DeNardis, Laura. 2009. *Protocol Politics: The Globalization of Internet Governance*. MIT Press.
- . 2014. *The Global War for Internet Governance*. Yale University Press.
- Department of Health, Education and Welfare. 1973. "Records, Computers and the Rights of Citizens." <https://epic.org/privacy/hew1973report/>.
- DePillis, Lydia. 2013. "There's a War in Cyberspace over Icons Vs. Text." *The New Republic*, January 17, 2013. <https://newrepublic.com/article/111970/app-terms-service-icons-or-text>.
- 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*.
- Dixon, Pam. 2007. "The Network Advertising Initiative: Failing at Consumer Protection and at Self-Regulation." World Privacy Forum. http://www.worldprivacyforum.org/wp-content/uploads/2007/11/WPF_NAI_report_Nov2_2007fs.pdf.
- "Do Not Track' Web System Stuck In Limbo." 2012. *All Things Considered*. <https://www.npr.org/2012/06/25/155727768/>.
- Doty, Nick. 2010. "Geolocation, Privacy and the Web." UC Berkeley TRUST seminar, September. <https://npdoty.name/slides/location-privacy-web2.pdf>.
- . 2013. "Because You Can Is Reason Enough to Do Something." *Bcc (blog)*. August 17, 2013. <http://bcc.npdoty.name/because-you-can-is-reason-enough-to-do-something>.

- . 2015a. “Reviewing for Privacy in Internet and Web Standard-Setting.” In *Security and Privacy Workshops (SPW), 2015 IEEE*, 185–92. IEEE. <https://npdoty.name/privacy-reviews/iwpe/>.
- . 2015b. “Interesting Questions I Heard from Students in Class: Standard-Setting and ‘Punting’ Decisions.” *Known.npdoty.name* (blog). February 24, 2015. <http://known.npdoty.name/2015/interesting-questions-i-heard-from-students-in-class-standard-setting-and>.
- . 2018. “Standard-Setting Process and Internet Privacy.” presented at the Protocol to the People? Protocol governance and power – from Bitcoin to the Border Gateway Protocol, Turing Institute, London, UK, March 16. <https://www.turing.ac.uk/events/protocol-people-protocol-governance-power-bitcoin-border-gateway-protocol/>.
- . 2019. “Mitigating Browser Fingerprinting in Web Specifications.” Interest Group Note. Privacy Interest Group (PING). World Wide Web Consortium. <https://www.w3.org/TR/2019/NOTE-fingerprinting-guidance-20190328/>.
- Doty, Nick, and Mohit Gupta. 2013. “Privacy Design Patterns and Anti-Patterns: Patterns Misapplied and Unintended Consequences.” In *A Turn for the Worse: Trustbusters for User Interfaces Workshop*. <http://cups.cs.cmu.edu/soups/2013/trustbusters.html>.
- 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/JHTL11I1-MulliganDoty.PDF>.
- Doty, Nick, Deirdre K Mulligan, and Erik Wilde. 2010. “Privacy Issues of the W3c Geolocation API.” *arXiv:1003.1775*, March. <http://arxiv.org/abs/1003.1775>.
- Doty, Nick, and Galen Panger. 2015. “Introducing Citizen Technologist, the Blog.” *CTSP Blog* (blog). September 9, 2015. <https://ctsp.berkeley.edu/introducing-citizen-technologist-the-blog/>.
- Doty, Nick, Heather West, Justin Brookman, Sean Harvey, and Erica Newland. 2019. “Tracking Compliance and Scope.” Working Group Note. Tracking Protection Working Group. World Wide Web Consortium. <https://www.w3.org/TR/tracking-compliance/>.
- 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>.
- Eaves, David, Ed Felten, Tara McGuinness, Deirdre K. Mulligan, and Jeremy Weinstein. 2020. “Defining Public Interest Technology.” *New America* (blog).

- January 22, 2020. <http://newamerica.org/pit/blog/defining-public-interest-technology/>.
- Ebbert, John. 2013. "IAB Vs Mozilla: Randall Rothenberg Takes The Gloves Off." *AdExchanger*, July 1, 2013. <https://www.adexchanger.com/online-advertising/iab-mozilla/>.
- Edelman, Gilad. 2020. "'Do Not Track' Is Back, and This Time It Might Work." *Wired*, October 7, 2020. <https://www.wired.com/story/global-privacy-control-launches-do-not-track-is-back/>.
- Einstein, Albert, and Bertrand Russell. 1955. "The Russell-Einstein Manifesto." Proceedings of the First Pugwash Conference on Science and World Affairs. <https://pugwash.org/1955/07/09/statement-manifesto/>.
- 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>.
- Farrell, S, and H Tschofenig. 2014. "Pervasive Monitoring Is an Attack." RFC 7258. RFC Editor. <http://tools.ietf.org/html/rfc7258>.
- Federal Trade Commission. 1998. "Privacy Online: A Report to Congress." <https://www.ftc.gov/reports/privacy-online-report-congress>.
- Federal Trade Commission, Bureau of Consumer Protection. 1983. "Standards and Certification: Final Staff Report." Washington, D.C. <https://catalog.hathitrust.org/Record/001535861>.
- Feng, Patrick. 2006. "Shaping Technical Standards: Where Are the Users?" In *Shaping Science and Technology Policy: The Next Generation of Research*, edited by David H. Guston and Daniel Sarewitz.
- Fisher, Roger, William L. Ury, and Bruce Patton. 2011. *Getting to Yes: Negotiating Agreement Without Giving in*. Penguin.
- Flanagan, M., D. Howe, and H. Nissenbaum. 2008. "Embodying Values in Technology: Theory and Practice." *Information Technology and Moral Philosophy*, 322–53.

- Ford, Paul. 2014. "The Group That Rules the Web." *The New Yorker*, November 20, 2014. <https://www.newyorker.com/tech/elements/group-rules-web>.
- Freedman, Tom, Jessica Roeder, Alexander Hart, Kyle Doran, and Kaye Sklar. 2016. "A Pivotal Moment: Developing a New Generation of Technologists for the Public Interest." Freedman Consulting. <http://tfreedmanconsulting.com/reports/a-pivotal-moment-developing-a-new-generation-of-technologists-for-the-public-interest/>.
- Freeman, Jody. 1997. "Collaborative Governance in the Administrative State." *UCLA L. Rev.* 45: 1.
- Freeman, Linton C. 1978. "Centrality in Social Networks Conceptual Clarification." *Social Networks* 1 (3): 215–39.
- 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.
- Gallie, W. B. 1956. "Essentially Contested Concepts." *Proceedings of the Aristotelian Society*, New Series, 56 (January): 167–98. <http://www.jstor.org/stable/4544562>.
- Geiger, RS, and David Ribes. 2011. "Trace Ethnography: Following Coordination Through Documentary Practices." In *System Sciences (HICSS), 2011 44th ...*, 0:1–10. Los Alamitos, CA, USA: IEEE Computer Society. <https://doi.org/10.1109/HICSS.2011.455>.
- Gillies, James, and R Cailliau. 2000. *How the Web was born: the story of the World Wide Web*. Oxford: Oxford University Press.
- Ginsberg, Allen. 1955. *Howl*. <https://www.poetryfoundation.org/poems/49303/howl>.
- Glaser, Barney G. 1978. *Theoretical Sensitivity: Advances in the Methodology of Grounded Theory*. Sociology 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>.
- Gürses, S., and J. M. del Alamo. 2016. "Privacy Engineering: Shaping an Emerging Field of Research and Practice." *IEEE Security Privacy* 14 (2): 40–46. <https://doi.org/10.1109/MSP.2016.37>.

- 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>.
- Hansen, Marit, John Morris, Alissa Cooper, Rhys Smith, Hannes Tschofenig, Jon Peterson, and Bernard Aboba. 2013. "Privacy Considerations for Internet Protocols." RFC 6973. Request for Comments. RFC Editor. <https://tools.ietf.org/html/rfc6973>.
- Harter, Philip J. 1982–1983. "Negotiating Regulations: A Cure for Malaise." *Georgetown Law Journal* 71: 1. <https://heionline.org/HOL/Page?handle=hein.journals/glj71&id=17&div=&collection=>.
- Hennink, Monique M, Bonnie N Kaiser, and Vincent C Marconi. 2017. "Code Saturation Versus Meaning Saturation: How Many Interviews Are Enough?" *Qualitative Health Research* 27 (4): 591–608.
- Hine, Christine. 2000. *Virtual Ethnography*. SAGE.
- "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/>.
- Hoofnagle, Chris Jay. 2016. *Federal Trade Commission Privacy Law and Policy*. Cambridge University Press.
- "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>.
- Jaffe, Jeff. 2018. "Diversity at W3c; Launch of TPAC Diversity Scholarship | W3c Blog." *W3c Blog* (blog). June 20, 2018. <https://www.w3.org/blog/2018/06/diversity-at-w3c-launch-of-tpac-diversity-scholarship/>.
- 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>.
- Knobel, Cory, and Geoffrey C. Bowker. 2011. "Values in Design." *Commun. ACM* 54 (7): 26–28. <https://doi.org/10.1145/1965724.1965735>.

- Koopman, Colin, and Deirdre K Mulligan. 2013. "Theorizing Privacy's Contestability: A Multi-Dimensional Analytic of Privacy."
- Kosack, Stephen, and Archon Fung. 2014. "Does Transparency Improve Governance?" *Annual Review of Political Science* 17: 65–87.
- Kostova, Blagovesta, Seda Gürses, and Carmela Troncoso. 2020. "Privacy Engineering Meets Software Engineering. On the Challenges of Engineering Privacy ByDesign." *arXiv:2007.08613 [Cs]*, July. <http://arxiv.org/abs/2007.08613>.
- Kumaraguru, Ponnurangam, and Lorrie Faith Cranor. 2005. "Privacy Indexes: A Survey of Westin's Studies."
- Latour, Bruno. 1987. *Science in Action : How to Follow Scientists and Engineers Through Society*. Harvard University Press. <http://www.amazon.com/dp/0674792904>.
- . 2007. *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford University Press, USA.
- 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>.
- Leon, Pedro, Blase Ur, Richard Shay, Yang Wang, Rebecca Balebako, and Lorrie Cranor. 2012. "Why Johnny Can'T Opt Out: A Usability Evaluation of Tools to Limit Online Behavioral Advertising." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 589–98. CHI '12. New York, NY, USA: ACM. <https://doi.org/10.1145/2207676.2207759>.
- Lessig, Lawrence. 1999. *Code and Other Laws of Cyberspace*. Basic Books. <http://books.google.com/books?id=011qLyT88XEC>.
- Lind, Edgar Allan, and Tom R. Tyler. 1988. *The Social Psychology of Procedural Justice*. Springer.
- Lipner, S. 2004. "The Trustworthy Computing Security Development Lifecycle." In *20th Annual Computer Security Applications Conference*, 2–13. <https://doi.org/10.1109/CSAC.2004.41>.
- Lofland, John, Lyn H. Lofland, David Snow, and Leon Anderson. 2006. *Analyzing Social Settings*. 4th ed. Wadsworth.
- Luegenbiehl, Heinz C., and Bill Puka. 1983. "Codes of Ethics and the Moral Education of Engineers [with Commentary]." *Business & Professional Ethics Journal* 2 (4): 41–66.

- Lynch, Michael. 2000. "Against Reflexivity as an Academic Virtue and Source of Privileged Knowledge." *Theory, Culture & Society* 17 (3): 26–54. <https://doi.org/10.1177/02632760022051202>.
- Lynch, William T., and Ronald Kline. 2000. "Engineering Practice and Engineering Ethics." *Science, Technology, & Human Values* 25 (2): 195–225.
- 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>.
- MacManus, Richard. 2003. "The Read/Write Web." *ReadWrite* (blog). April 20, 2003. https://readwrite.com/2003/04/19/the_readwrite_w/.
- Manjoo, Farhad. 2013. "The Dumb Argument That 3-D Printers Will Make Gun Control Futile." *Slate Magazine*. May 8, 2013. <https://slate.com/technology/2013/05/3-d-printed-gun-yes-it-will-be-possible-to-make-weapons-with-3-d-printers-no-that-doesnt-make-gun-control-futile.html>.
- Marcus, GE. 1983. *Elites: Ethnographic Issues*. Albuquerque: SAR Press. <http://www.amazon.com/Elites-Ethnographic-American-Research-Advanced/dp/193469133X>.
- Massey, Doreen B. 1994. *Space, Place, and Gender*. U of Minnesota Press.
- 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>.
- Matias, J. Nathan. 2014. "How to Ethically and Responsibly Identify Gender in Large Datasets." *MediaShift*. November 21, 2014. <http://mediashift.org/2014/11/how-to-ethically-and-responsibly-identify-gender-in-large-datasets/>.
- "Mechanism Not Policy." 2005, July. <http://c2.com/cgi/wiki?MechanismNotPolicy>.
- "Memorandum of Understanding Between W3c and WHATWG." 2019. May 28, 2019. <https://www.w3.org/2019/04/WHATWG-W3C-MOU.html>.
- Molla, Rani, and Renee Lightner. 2016. "Diversity in Tech Companies." *Wall Street Journal*. April 10, 2016. <http://graphics.wsj.com/diversity-in-tech-companies/>.
- 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/>.

- Morozov, Evgeny. 2013. *To Save Everything, Click Here: The Folly of Technological Solutionism*. Public Affairs.
- Mulligan, Deirdre K., and Nick Doty. 2016. "Design Wars: The FBI, Apple and Hundreds of Millions of Phones." *Citizen Technologist: The CTSP Blog* (blog). March 3, 2016. <https://ctsp.berkeley.edu/design-wars-fbi-apple/>.
- Mulligan, Deirdre K., Colin Koopman, and Nick Doty. 2016. "Privacy Is an Essentially Contested Concept: A Multi-Dimensional Analytic for Mapping Privacy." *Phil. Trans. R. Soc. A* 374 (2083): 20160118. <https://doi.org/10.1098/rsta.2016.0118>.
- Mulligan, Deirdre K, and Helen Nissenbaum. 2020. "Handoffs." *In Progress*.
- Nader, Laura. 1972. "Up the Anthropologist: Perspectives Gained from Studying Up."
- Nissenbaum, Helen. 1998. "Values in the Design of Computer Systems." *Computers and Society* 28 (1): 38–39. <http://www.nyu.edu/projects/nissenbaum/papers/society.pdf>.
- . 2004. "Privacy as Contextual Integrity." *Washington Law Review* 79 (1): 101–39. http://heinonlinebackup.com/hol-cgi-bin/get_pdf.cgi?handle=hein.journals/washlr79§ion=16.
- NIST. 2014. "NIST Privacy Engineering Objectives and Risk Model Discussion Draft." April. http://csrc.nist.gov/projects/privacy%7B-%7Dengineering/nist_privacy_engr_objectives_risk_model_discussion_draft.pdf.
- . 2015. "Privacy Risk Management for Federal Information Systems." http://csrc.nist.gov/publications/drafts/nistir-8062/nistir-8062_draft.pdf.
- 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>.
- O'Neill, Mike. 2018. "Do Not Track and the GDPR." *W3c Blog* (blog). June 11, 2018. <https://www.w3.org/blog/2018/06/do-not-track-and-the-gdpr/>.
- "Obligation." 2018. Order of The Engineer. 2018. http://www.order-of-the-engineer.org/?page_id=6.
- Oever, Niels ten, and Davide Beraldo. 2018. "Routes to Rights: Internet Architecture and Values in Times of Ossification and Commercialization." *XRDS: Crossroads, The ACM Magazine for Students* 24 (4): 28–31. <https://doi.org/10.1145/3220561>.
- Okumura, Kaori, Yoshiaki Shiraishi, and Akira Iwata. 2013. "Survey on Sense of Security for Registering Privacy Information to Return Refugee Supporting System." In *Symposium on Usable Privacy and Security (SOUPS)*. <https://cups>

- .cs.cmu.edu/soups/2013/trustbusters2013/Sense_of_Security_Refugee_Supporting_System_0kumura.pdf.
- Organization for Economic Cooperation and Development. 1980. "Guidelines on the Protection of Privacy and Transborder Flow of Personal Data." <http://www.oecd.org/internet/ieconomy/oecdguidelinesontheProtectionofPrivacyandTransborderFlowsofPersonalData.htm>.
- Ortega y Gasset, José, and John William Miller. 1962. *History as a System and Other Essays Toward a Philosophy of History*. Translated by Helene Weyl. New Ed edition. New York: W. W. Norton & Company.
- 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=snug08KeC2EC>.
- 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>.
- Quine, W. V. 1951. "Two Dogmas of Empiricism." *The Philosophical Review* 60 (1): 20–43. <https://doi.org/10.2307/2181906>.
- 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>.
- Sandelowski, M. 1986. "The problem of rigor in qualitative research." *Advances in nursing science* 8 (3): 27–37. <http://www.ncbi.nlm.nih.gov/pubmed/3083765>.
- Savage, Charlie. 2013. "U.S. Weighs Wide Overhaul of Wiretap Laws." *The New York Times*, May 8, 2013, sec. U.S. <https://www.nytimes.com/2013/05/08/us/politics/obama-may-back-fbi-plan-to-wiretap-web-users.html>.
- Saxenian, AnnaLee. 1996. *Regional Advantage*. Harvard University Press.
- Schechter, Emily. 2016. *Moving Towards a More Secure Web*. <https://security.googleblog.com/2016/09/moving-towards-more-secure-web.html>.
- Sebenius, James K. 1983/ed. "Negotiation Arithmetic: Adding and Subtracting Issues and Parties." *International Organization* 37 (2): 281–316. <https://doi.org/10.1017/S002081830003438X>.

- “Self-Regulation and Privacy Online,’ FTC Report to Congress.” 1999. Federal Trade Commission. July 13, 1999. <https://www.ftc.gov/news-events/press-releases/1999/07/self-regulation-and-privacy-online-ftc-report-congress>.
- “Self-Review Questionnaire: Security and Privacy.” 2020. W3c Technical Architecture Group. World Wide Web Consortium. <https://w3ctag.github.io/security-questionnaire/>.
- Sennett, Richard. 2008. *The Craftsman*. London: Allen Lane.
- Siegman, Tzviya, An Qui Li, and Ada Rose Cannon. 2020. “Positive Work Environment at W3c: Code of Ethics and Professional Conduct.” World Wide Web Consortium. <https://www.w3.org/Consortium/cepc/cepc-20200716/>.
- 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>.
- Snyder, Pete. 2019. “Privacy Anti-Patterns in Standards.” *W3c Blog* (blog). June 12, 2019. <https://www.w3.org/blog/2019/06/privacy-anti-patterns-in-standards/>.
- Solove, DJ. 2006. “A Taxonomy of Privacy.” *University of Pennsylvania Law Review*, no. 477: 477–560. <http://www.jstor.org/stable/10.2307/40041279>.
- 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>.
- Stark, Luke, and Anna Lauren Hoffmann. 2019. “Data Is the New What? Popular Metaphors & Professional Ethics in Emerging Data Culture.” *Journal of Cultural Analytics*. <https://doi.org/10.22148/16.036>.
- Strauss, Anselm, and Juliet Corbin. 1990. *Basics of Qualitative Research*. Sage publications.
- Sunshine, Joshua, Serge Egelman, Hazim Almuhiemedi, Neha Atri, and Lorrie Faith Cranor. 2009. “Crying Wolf: An Empirical Study of SSL Warning Effectiveness.” In *USENIX Security Symposium*, 399–416.
- Sunstein, Cass R. 2014. “From Technocrat to Democrat.” *Harvard Law Review* 128 (1): 488–97.
- 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 ‘Do Not Track’ Setting Doesn’t Stop You from Being Tracked.” 2019. *DuckDuckGo Blog* (blog). February 5, 2019. <https://spreadprivacy.com/do-not-track/>.
- “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/>.
- Thomson, Martin. 2014. “A Statement.” Internet-Draft draft-thomson-perpass-statement-01. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/html/draft-thomson-perpass-statement-01>.
- Tufekci, Zeynep. 2016. “The Real Bias Built In at Facebook.” *The New York Times*, May 19, 2016, sec. Opinion. <https://www.nytimes.com/2016/05/19/opinion/the-real-bias-built-in-at-facebook.html>.
- 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.
- Vance, Ashlee. 2011. “This Tech Bubble Is Different.” *Bloomberg Businessweek*, April 14, 2011. <https://www.bloomberg.com/news/articles/2011-04-14/this-tech-bubble-is-different>.
- Wallace, David Foster. 2005. “This Is Water.”
- Warren, Samuel D., and Louis D. Brandeis. 1890. “The Right to Privacy.” *Harvard Law Review*, 193–220.
- Watt, Diane. 2007. “On Becoming a Qualitative Researcher: The Value of Reflexivity.” *Qualitative Report* 12 (1): 82–101.
- 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>.
- Westin, A. 2001. “Privacy on & Off the Internet: What Consumers Want.” Technical report, Tech. Report for Privacy & American Business. Hackensack, NJ: Privacy & American Business.
- Westin, A.F. 1967. *Privacy and Freedom*. New York: Atheneum.
- Winner, Langdon. 1980. “Do Artifacts Have Politics?” *Daedalus* 109 (1): 121–36. <http://www.jstor.org/stable/20024652>.

Wong, Richmond Y., and Deirdre K. Mulligan. 2019. "Bringing Design to the Privacy Table: Broadening 'Design' in 'Privacy by Design' Through the Lens of HCI." In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–17.

Appendices

Appendix: Interview Guide

Interview Guide – Privacy in Technical Standard-Setting

Your role and privacy

- What is your role at X?
 - (only if it comes up) what background led you to this role?
- How do you think about privacy in your work?
 - Can you tell me about a time when privacy came up in a product discussion?
 - When does it get discussed among your colleagues?
- What types of privacy threats do you consider?
 - Do you distinguish between privacy and security?
 - privacy from-whom, attacker, collection/aggregation/use

Personal views of privacy

- Do you consider yourself a private person? When was a time you remember worrying about your own privacy? [you can tell me in general terms or of course leave out details if you like] Why was it a concern?
 - Online? Offline?
 - [pick up on words used and ask more]
 - [if only one is mentioned] Do you worry about privacy from the government? From your family and friends? From corporations? From strangers?
- When you're designing software or debating standards, how do you imagine that users of the Web think about privacy?
 - Do you agree with them? Why or why not?

Technical standards and privacy

- And how did you first get involved in technical standards?
 - How does your role interact with standards work? (What other roles in your organization are involved?)
- Can you tell me about a time you remember that privacy came up in a standards group you were part of?
 - How did it play out?
 - What was your role?
 - Was there a debate? What were the sides? What was the outcome?
 - ... Can you remember a time where you supported a different side? Or where there was a different outcome?
- How did you make your decision about the privacy debate?
 - Are you representing your company? The user? The best technical solution?
 - Did this debate belong in standardization? Why or why not?
 - What is the role/purpose of standards here? [interoperability; designing better technology; consistency; fundamental values?]
- Can you tell me about a time you remember that legal considerations came up in a standards group?
 - Does this happen frequently? What about within your company?
 - What about accessibility? Security?
- Are you satisfied with how privacy has been handled in standards discussions? Why or why not? What counts as successful or not?
 - What would you do differently? What would improve it? How would that have helped?

optional: Do Not Track process *For participants or non-participant stakeholders in the W3C Tracking Protection Working Group.*

- How would you summarize your experience with (or impressions of) the Do Not Track standardization process?
- How did you see the role of W3C standardization in the larger privacy debate?

- What were your goals for the process?
 - to what extent were they met? what counts as successful for you?
 - what would you have preferred was done differently? How would that have helped?
- Do you think the process was fair? Why or why not?
- Do you think the process produced a good outcome? Why or why not?
- Do you recall a particular debate in the Working Group that involved you?
 - how did it play out?
 - what was your role?
- How did the involvement of some other participants affect the process?
 - [suggest categories of advertisers, browser vendors, advocates, regulators, etc.]
 - [prompt about press/publicity if it doesn't come up]
 - how have views of participants changed from before the process started?
- How do you expect your working relationship will continue with the other participants in the group?

Westin index Would you strongly agree, agree, disagree or strongly disagree with the following?

1. Consumers have lost all control over how personal information is collected and used by companies.
2. Most businesses handle the personal information they collect about consumers in a proper and confidential way.
3. Existing laws and organizational practices provide a reasonable level of protection for consumer privacy today.

Wrap-up

- Who else should I talk to? Were there particular people you recall holding particular positions in privacy debates?

Appendix: Privacy-as-x

A mostly exhaustive list of coding of concepts of privacy (privacy-as-x) based on interviewees' volunteered language around privacy:

- anonymity
- autonomy
- protection from a chilling effect
- compliance
- confidentiality or secrecy
- consent and control / informational self-determination
- consumer trust
- contextual integrity
- not being creepy
- data handling, data governance, data protection
- disconnection (ability to be offline, free from distracting tech)
- discretion, etiquette, relevance (not “too much information”)
- protection from embarrassment
- awareness/limitation of inferences
- freedom from intrusion (annoyance / harassment)
- redress / correction / deletion
- risk mitigation
- not surprising the user
- transparency