

December 2, 2020

Who participates and why it matters

Nick Doty

UC Berkeley, School of Information

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

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.

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 convening¹), 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.²

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 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.

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?”³ 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.

¹See the stages of success in standard-setting process.

²See *Internet Standard-Setting and Multistakeholder Governance* in the section “Legitimacy and interoperability” and in *Doty and Mulligan* (2013).

³In this case, quoting and citing feminist scholars who are critical of Habermasian discourse theory.

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,⁴ 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 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.⁵ 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

⁴See Doty and Mulligan (2013) and [Internet Standard-Setting and Multistakeholder Governance](#).

⁵Ortega y Gasset and Miller (1962), as detailed previously in “Separation vs. Integration” in [The Ethics of Engineering](#).

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,⁶ diversity of participation is identified as especially important for questions of policy or ethical values.

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

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

⁶Namely, 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](#).)

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 BigBang⁷ to crawl, parse and consolidate mailing list archives. The automated analysis here makes use of Gender Detector,⁸ 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.

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.

For a corpus of all active W3C Working Groups and Interest Groups as of 2017,⁹ we can estimate the fraction of male and female gender among participants who sent at least one message to those mailing lists. Those results are presented in the table, Figure 1.

As we would expect, most groups seem to have mostly participants inferred to be men. While many participants' genders can't effectively be estimated this way, nonetheless the average fraction of participants that are identified as men in one of these working or interest groups is over half, while on average only 8% are identified as women. These averages may provide a useful baseline and a point of comparison for future automated or semi-automated estimates. Diversity reports published by several major tech companies, all of whom participate to some degree in Internet and Web standard-setting, provide one point of comparison. In data from 2015, the percentage of technology jobs held by women ranged from 13% at Twitter to 24% at eBay (Molla and Lightner 2016). And while larger W3C surveys are not currently available, there are reports on the demographic breakdown of some leadership groups: as of 2018, the Technical Architecture Group was made up of 10% women, although the Advisory Board was closer to parity (Jaffe 2018).

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: ac-

⁷<https://github.com/dataactive/bigbang> and see the [Methods chapter](#) for more description.

⁸<https://github.com/malev/gender-detector>

⁹While we could also run this on more recent or larger datasets, this one is of interest for my purposes as it includes the Tracking Protection Working Group and the other focused groups at W3C only that were active around the same time.

will be dominant.

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,¹¹ 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.

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

¹¹A *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.

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.”¹² 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.

¹²This 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.

Relevance of stakeholder group analysis Is the level of granularity really so important?¹³ 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.¹⁴ For example, assuming 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

¹³It’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](#)).

¹⁴Again, see [methods chapter](#).

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.

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.

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.”¹⁵ 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.

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.¹⁶

¹⁵This 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.

¹⁶This is discussed further in [How participants see privacy](#).

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.

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.¹⁷

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 (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

¹⁷This 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 (2015).

leadership roles. The highest ranked woman is Alissa Cooper, current Chair of the IETF.¹⁸

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:

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.¹⁹

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

¹⁸ Annotated 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>

¹⁹ Emerson et al. (2009), as described in Doty and Mulligan (2013).

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 uncommon²⁰ 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.

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](#).

²⁰Indeed, it sometimes surprised me, despite the hypothesis in our previous work on this topic.

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.

References

- Doty, Nick. 2015. "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/>.
- Doty, Nick, and Deirdre K. Mulligan. 2013. "Internet Multistakeholder Processes and Techno-Policy Standards: Initial Reflections on Privacy at the World Wide Web Consortium." *Journal on Telecommunications and High Technology Law* 11. http://www.jthtl.org/content/articles/V11I1/JTHTLv11i1_MulliganDoty.PDF.
- 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>.
- 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.
- Freeman, Linton C. 1978. "Centrality in Social Networks Conceptual Clarification." *Social Networks* 1 (3): 215–39.
- Froomkin, A. Michael. 2003. "Habermas@Discourse. Net: Toward a Critical Theory of Cyberspace." *Harvard Law Review* 116 (3): 749–873. <https://doi.org/10.2307/1342583>.
- 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/>.
- 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/>.
- Molla, Rani, and Renee Lightner. 2016. "Diversity in Tech Companies." *Wall Street Journal*. April 10, 2016. <http://graphics.wsj.com/diversity-in-tech-companies/>.
- 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.
- Quine, W. V. 1951. "Two Dogmas of Empiricism." *The Philosophical Review* 60 (1): 20–43. <https://doi.org/10.2307/2181906>.
- "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/>.