

December 15, 2020

How standard-setting accommodates, succeeds and fails

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](#).*

How standard-setting accommodates, succeeds and fails

In considering the future of multistakeholderism for tech policy, I asked:¹ 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,

¹See [Internet Standard-Setting and Multistakeholderism](#).

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.

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 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 sides² 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

²See [Stakeholder groups: counting sides](#) in the section on participation.

opposition to Do Not Track and the W3C process, interspersed into a technical discussion of unlinkability.³ 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.

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.

³Minutes from October 2012 and press release re: open letter from DAA to W3C leadership.

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.

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.

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;⁴ those efforts have also faced pushback which has typically demonstrated the presence of discouraging and antagonistic behavior and the need for more welcoming environments.

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.

⁴Positive Work Environment statement of principles [appears to date to June 2007](#) and a more formal version was published this year [Siegman, Li, and Cannon \(2020\)](#).

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 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.⁵ 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

⁵Interview was in 2018, but this is referring primarily to 2011-2013.

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.

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.⁶

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

⁶Several heuristics might still be partial explanations, though, including also gender, professional background and cultural attitudes.

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?

“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 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.⁷

⁷See “Standard-setting organizations as social networks” in the section on participation for more on this topic.

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? 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.

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

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.

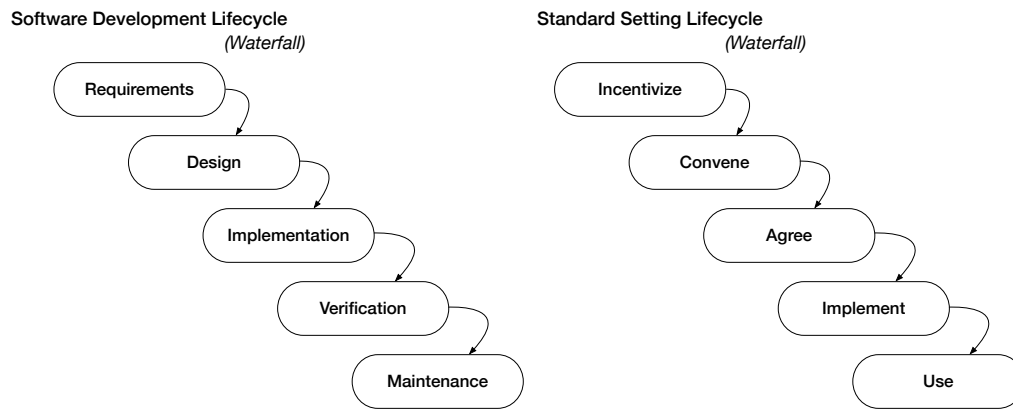


Figure 1: The traditional waterfall model of the software development lifecycle, with a somewhat analogous set of steps for developing a technical standard.

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 now⁸ 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.

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.⁹ Convening has also been a key tool of FTC and other agencies in pursuing new governance approaches rather than more traditional, formal rule-making.¹⁰

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

⁸Circa 2020.

⁹As described in the chapter on [Internet Standard-Setting and Multistakeholder Governance](#), citing [Bruant \(2013\)](#) and a description of W3C.

¹⁰See, for example, [Cohen \(2012\)](#) and [Bamberger and Mulligan \(2010\)](#).

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.

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 perspec-

tives 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 plusses 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.

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,¹¹ (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.

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,¹² 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.

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.”

¹¹See Sebenius (1983/ed) for one relevant description.

¹²See *Individuals vs organizations*.

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?

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.

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.

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)¹³. 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 adop-

¹³Numbers 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).

tion 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 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.¹⁴

¹⁴Unofficial draft: <https://globalprivacycontrol.github.io/gpc-spec/>

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.

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.

References

- Bamberger, Kenneth A., and Deirdre K. Mulligan. 2010. "Privacy on the Books and on the Ground." *Stan. L. Rev.* 63: 247.
- Bray, Tim. 2012. "On the Deadness of OAuth 2." *Ongoing* (blog). July 28, 2012. <https://www.tbray.org/ongoing/When/201x/2012/07/28/0auth2-dead>.
- 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/>.
- Cohen, Julie E. 2012. "What Privacy Is For." *Harv. L. Rev.* 126: 1904.
- 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/>.
- Fisher, Roger, William L. Ury, and Bruce Patton. 2011. *Getting to Yes: Negotiating Agreement Without Giving in*. Penguin.
- 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/>.

- 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>.
- 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/>.
- "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/>.